

Reconstructing Input Models in Stochastic Simulation

Aleksandrina Goeva

Department of Mathematics and Statistics, Boston University, Boston, MA 02215, USA. agoeva@bu.edu

Henry Lam

Department of Industrial and Operations Engineering, University of Michigan, Ann Arbor, MI 48109, USA. khlam@umich.edu

Huajie Qian

Department of Mathematics, University of Michigan, Ann Arbor, MI 48109, USA. hqian@umich.edu

Bo Zhang

IBM Research, Yorktown Heights, NY 10598, USA. zhangbo@us.ibm.com

We investigate the problem of nonparametrically calibrating the input model in stochastic simulation, given only the availability of output data. While studies on simulation input uncertainty have focused on the use of input data, our inverse model calibration problem arises in various situations where resource and operational limitations hinder the direct observability of input data. We propose a moment-based, maximum entropy framework to infer the input model, by matching statistics between the simulation output and the real-world output. To bypass the methodological difficulties from the stochastic constraints in our formulation, we propose a stochastic quadratic penalty method that converts the problem into a sequence of least-square problems, where each element in the sequence can be solved by efficient stochastic approximation algorithms. We analyze the statistical properties of our method and demonstrate its ability to recover input models with numerical experiments.

Key words: model calibration; maximum entropy; moment methods; nonparametric; simulation optimization

1. Introduction

Stochastic simulation takes in input models and generate random outputs for subsequent performance analyses. The accuracy of these input model assumptions is critical to the analyses' credibility. In the conventional premise in studying stochastic simulation, these input models are conferred either through physical implication or expert opinion, or observable via input data. In this paper, we ask a converse question: Given only *output* data from a stochastic system, can one infer about the input model?

The main motivation of asking this question is that, in many situations, a simulation modeler plainly may not have the luxury of direct data or knowledge about the input. The only way to gain such knowledge could be data from other sources that are at the output level. For instance,

one of the authors has experienced such complication when building a simulation model for a contract fulfillment center, where service agents work on a variety of processing tasks and, despite the abundant transaction data stored in the center’s IT system, there is no record on the start, completion, or service times spent by each agent on each particular task. Similarly, in clinic operations, patients often receive service in multiple phases such as initial checkup, medical tests and doctor’s consultation. Patients’ check-in and check-out times could be accurately noted, but the “service” times provided by the medical staff could very well be unrecorded. Clearly, these service time distributions are needed to build a simulation model, if an analyst wants to use the model for sensitivity or system optimization purpose.

The problem of inferring an input model from output data is sometimes known as *model calibration*. In the simulation literature, this is often treated as a refinement process that occurs together with iterative comparisons between simulation reports and real-world output data (a task known as *model validation*; Sargent (2005), Kleijnen (1995)). If simulation reports differ significantly from output data, the simulation model is re-calibrated (which can involve both the input distributions and system specifications), re-compared, and the process is iterated. Suggested approaches to compare simulation with real-world data is to conduct statistical tests such as two-sample mean-difference tests (Balci and Sargent (1982)) and the Schruben-Turing test (Schruben (1980)). Beyond that, inferring input from output seems to be an important problem that has not been largely discussed in the stochastic simulation literature (Nelson (2016)). Thus, on a high level, one contribution in this paper is to study the first systematic framework for the input model calibration problem.

The setting we consider can be briefly described as follows. We assume an input model is missing and make no particular assumptions on the form of its probability distribution. We assume, however, that certain output random variable from a well-specified system is observable with some data. Our task is to nonparametrically infer the input distribution. A key insight we use that distinguishes this problem with model calibration in other literature (e.g., computer experiments) is the intrinsic probabilistic structure in the system. Namely, the input and the output in stochastic simulation are represented as probability distributions, or in other words, the relation that links the observed and the to-be-calibrated objects is a (simulable) map between the spaces of distributions. Our calibration method is, and should be, designed to take such a relation into account.

Specifically, we use a moment-based approach to calibrate the input model. We match the moments of a sequence of statistics collected from the simulation and real-world data at the output level. Unless we can afford to match an infinite number of such moments, however, this scheme in general can lead to many candidate input models. To tackle this issue, we propose entropy maximization as a criterion to identify the best model. This proposal is based on the interpretation of

the maximum entropy (ME) distribution as the conditional distribution given all prior information, hence the most “natural” distribution without any further knowledge (Van Campenhout and Cover (1981)), and also its use in density estimation (Barron and Sheu (1991)). We shall offer some justification in the consistency of our moment-matching approach, while at the same time we shall also reveal some finite-sample limitations. Roughly speaking, a “natural” model such as unimodal and smooth distribution is amenable with our approach, but a “complicated” distribution is more difficult to infer. We shall provide some numerical results to substantiate these observations.

Beyond proposing and justifying our formulation, our another key contribution is algorithmic. In particular, we study a simulation optimization procedure that targets at our ME moment-matching problem. A challenging characteristic of this problem is that the moment statistics are at the output level, which is generally a nonlinear function of the input distribution (as a simple example, think about the mean waiting time of a queue as a function of the service time distribution). Thus the constraints in the optimization formulation are stochastic and non-convex. There is essentially no literature on this class of simulation optimization problems. As our key algorithmic contribution, we propose and analyze a *stochastic quadratic penalty method*. The idea is to reformulate the optimization into a sequence of sum-of-square optimizations with stochastic objectives and deterministic convex constraints, whereby each element of the sequence can be solved efficiently using mirror descent stochastic approximation (MDSA) (Nemirovski et al. (2009)), a constrained SA algorithm that is especially suitable for our setting. We analyze the details of our reformulation. We also develop the gradient estimation procedure and analyze the convergence of our MDSA. Note that, for a general simulation model, the ME formulation and our reformulation only guarantees a local optimum because of non-convexity. In practice, this issue can be alleviated with the use of multi-start procedures. We also propose some visualization tool to validate our method.

The remainder of the paper is organized as follows. Section 2 reviews the related literature. Section 3 introduces the problem setting and explains our optimization formulation. Section 4 presents and analyzes our algorithm. Section 5 reports numerical results. Section 6 concludes and discusses future work. The Appendix contains some auxiliary theorems. The Supplementary Materials contains the details of an additional variant of our method.

2. Related literature

We organize the literature review in two aspects, one related to the model calibration problem, and one related to our optimization approach.

2.1. Literature Related to Our Problem Setting

Input modeling and uncertainty in the stochastic simulation focuses mostly on the input level. Barton (2012) and Song et al. (2014) review some major methods in quantifying the statistical errors from finite input data. These methods include the delta or two-point method

(Cheng and Holland (1998, 2004)), Bayesian methodology and model averaging (Chick (2001), Zouaoui and Wilson (2004)) and resampling methods (Barton and Schruben (2001), Barton et al. (2013)). Our problem is more related to model calibration. In the simulation literature, this is often considered together with model validation (Sargent (2005), Kleijnen (1995)). Conventional approaches compare simulation data with real-world historical output data according to statistical or Turing tests (Balci and Sargent (1982), Schruben (1980)), conduct re-calibration, and repeat the process until the data are successfully validated (Banks et al. (2009), Kelton and Law (2000)).

The model calibration problem is also known as the *inverse problem* (Tarantola (2005)) in other literatures. It generally refers to the identification of parameters or functions that can only be inferred from transformed outputs. In signal processing, the linear inverse problem (e.g., Csiszár (1991), Donoho et al. (1992)) reconstructs signals from measurements of linear transformations. Common approaches consist of least-square minimization and the use of penalty such as the entropy, which also provide justifications to our main formulation. In computer experiments (Santner et al. (2013)), surrogate models built on complex physical laws require the calibration of physical parameters. Such models have wide scientific applications such as weather prediction, oceanography, nuclear physics, and acoustics (e.g., Wunsch (1996), Shirangi (2014)). Bayesian and Gaussian process methodologies are commonly used (e.g., Kennedy and O’Hagan (2001), Currin et al. (1991)). We point out that Bayesian methods could be a potential alternative to the approach considered in this paper, but because of the nature of discrete-event systems, one might need to resort to sophisticated techniques such as approximate Bayesian computation (Marjoram et al. (2003)).

Also related to our work is the body of research on inference problems in the context of queueing systems. The first stream, similar to our paper, aims at inferring the constituent probability distributions of a queueing model based on its output data, e.g., queue length or waiting time data, collected either continuously or at discrete time points. This stream of papers focuses on systems whose structures allow closed-form analyses or are amenable to analytic approximations via, for instance, the diffusion limit. The majority of them assume that the inferred distribution(s) comes from a parametric family and use maximum likelihood estimators (Basawa et al. (1996), Pickands III and Stine (1997), Basawa et al. (2008), Fearnhead (2004), Wang et al. (2006), Ross et al. (2007), Heckmüller and Wolfinger (2009), Whitt (2012)). Others work on nonparametric inference by exploiting specific queueing system structures (Bingham and Pitts (1999), Hall and Park (2004), Moulines et al. (2007), Feng et al. (2014)). A related stream of literature studies point process approximation (see Section 4.7 of Cooper (1972), Whitt (1981, 1982), and the references therein), based on a parametric approach and is motivated from traffic pattern modeling in communication networks. Finally, there are also a number of studies inspired by the “queue inference engine” by Larson (1990). But, instead of inferring the input models, many of

these studies use transaction data to estimate the performance of a queueing system directly and hence do not take on the form of an inverse problem (see Mandelbaum and Zeltyn (1998) for a good survey of the earlier literature and Frey and Kaplan (2010) and its references for more recent progress). Several papers estimate both the queueing operational performance and the constituent input models (e.g., Daley and Servi (1998), Kim and Park (2008), Park et al. (2011)), and can be considered to belong to both this stream and the aforementioned first stream of literature.

2.2. Literature Related to Our Methodology

Our formulation uses the widely used notion of maximum entropy (ME). In information theory, entropy can be viewed as the amount of intrinsic randomness (Cover and Thomas (1991)). Csiszár (1991) studied axiomatic properties of ME (or more generally the I -divergence class). Donoho et al. (1992) used ME to recover sparse signals. Barron and Sheu (1991) studied the use of moment-constrained ME in nonparametric density estimation and analyzed the convergence rate in terms of Kullback-Leibler (KL) divergence. Lindley (1956), DeGroot (1962), Box and Hill (1967), Bernardo (1979), Chick and Ng (2002) studied the use of entropy criterion to maximize information and identify important parameters in experiments. In finance, Avellaneda et al. (2001) studied ME calibration of risk-neutral measures from derivative prices. This technique, known as the weighted Monte Carlo, has also been studied as a variance reduction technique (Glasserman and Yu (2005)).

Our optimization reformulation is inspired by the quadratic penalty method (Bertsekas (1999)), which is a deterministic nonlinear programming technique that reformulates the constraints as squared penalty and sequentially tunes the penalty parameter to approach optimality. Our algorithm to solve the sequence of optimizations in the reformulation utilizes MDSA proposed by Nemirovski et al. (2009). Nemirovski et al. (2009) analyzed convergence guarantees on convex programs with stochastic objectives. Ghadimi and Lan (2013) investigated related methods for nonconvex programs, and Ghadimi and Lan (2015) and Dang and Lan (2015) studied generalizations incorporating accelerated gradient and coordinate decomposition. The particular scheme of MDSA we consider uses entropic penalty, and is known as the entropic descent algorithm (Beck and Teboulle (2003)).

3. Setting and Formulation

We assume a discrete probability distribution $\mathbf{p} = (p_1, \dots, p_n)$ for the input model, on the support set $\mathcal{S} = \{z_1, \dots, z_n\}$, where the support size n can be potentially large. We let $\mathbf{X} = (X_1, \dots, X_\tau)$, where $X_t \in \mathbb{S}$, be an i.i.d. sequence of input variates each distributed under \mathbf{p} over a random horizon τ . We denote the function $h(\cdot) \in \mathbb{R}$ as the system logic from the input sequence \mathbf{X} to the output $h(\mathbf{X})$. We assume that h is completely specified and is computable, even though it may not be writable in closed-form, i.e. we can evaluate the output given \mathbf{X} . For example, \mathbf{X} can

denote the sequence of interarrival or service times for the customers in a queue, and $h(\mathbf{X})$ is the average queue length until the first idle time. Note that we can work in a more general framework where h depends on both \mathbf{X} and other independent input sequences, say \mathbf{Y} , that possess known or observable distributions. In other words, we can have $h(\mathbf{X}, \mathbf{Y})$ as the output. Our subsequent discussion can be trivially extended to this case, and hence we will suppress these auxiliary input sequences throughout our exposition.

Consider the situation that only $h(\mathbf{X})$ can be observed via data. Let y_1, \dots, y_N be N observations of $h(\mathbf{X})$. Our task is to calibrate \mathbf{p} .

We match the moment-based statistics of the simulation output and the empirical output data. More precisely, let $\phi_j(\cdot) : \mathbb{R} \rightarrow \mathbb{R}, j = 1, \dots, m$ be m moment functions specified by the modeler. Natural examples of ϕ_j include polynomials $\phi_j(y) = y^j$, and quantile-based functions $\phi_j(y) = I(y \leq c_j)$ for given values c_j , where I denotes the indicator function. We want to find \mathbf{p} such that

$$E_{\mathbf{p}}[\phi_j(h(\mathbf{X}))] = \hat{\mu}_j \quad \text{for } j = 1, \dots, m \quad (1)$$

where $E_{\mathbf{p}}[\cdot]$ denotes the expectation with respect to the i.i.d. input process \mathbf{X} each distributed as \mathbf{p} (in other words, the product measure $\mathbf{p} \times \mathbf{p} \times \dots$) and the random time τ . $\hat{\mu}_j$ is the empirical moment

$$\hat{\mu}_j = \frac{1}{N} \sum_{r=1}^N \phi_j(y_r), \quad j = 1, \dots, m \quad (2)$$

Note that, when the support size n is larger than m , there are typically more than one \mathbf{p} that satisfies (1). Moreover, depending on what the function h is, there may be a fundamental barrier in fully recovering \mathbf{p} even if there is a full knowledge on the distribution of $h(\mathbf{X})$. For example, when h is identically equal to a constant, any \mathbf{p} will give a perfect fit. The phenomenon of being unable to fully recover \mathbf{p} is generally known as non-identifiability in the inverse modeling literature (e.g. Tarantola (2005)).

Our approach to reduce the number of distributions \mathbf{p} that satisfy (1) is to maximize entropy, namely

$$\begin{aligned} \max_{\mathbf{p} \in \mathcal{P}} \quad & R(\mathbf{p}) = - \sum_{i=1}^n p_i \log p_i \\ \text{subject to} \quad & E_{\mathbf{p}}[\phi_j(h(\mathbf{X}))] = \hat{\mu}_j \quad \text{for } j = 1, \dots, m \end{aligned} \quad (3)$$

where $\mathcal{P} = \{\mathbf{p} : \sum_{i=1}^n p_i = 1, p_i \geq 0 \text{ for } i = 1, \dots, n\}$ is the probability simplex on $\mathcal{S} = \{z_1, \dots, z_n\}$. $R(\mathbf{p})$ denotes the entropy of \mathbf{p} . The decision variable in (3) is the unknown input model \mathbf{p} .

3.1. Why Moment Matching?

Moment estimators are ubiquitous in parametric estimation (e.g., Hall (2005)). In our framework, which is closer to nonparametric, the moment-matching is based on two beliefs: 1) Given sufficiently enough moments, one can recover the distribution of the output to high accuracy; 2) There is a

one-to-one map from the input distribution to the output distribution. These two beliefs together would lead to the recovery of the input model to high accuracy given enough output data and moments matched. The following presents this logic more rigorously:

LEMMA 1. Denote $F_Y : \mathbb{R} \rightarrow [0, 1]$ as the distribution function of the output $Y = h(\mathbf{X})$. We write $F_Y(\mathbf{p})$ to highlight the dependence of F_Y in terms of the input distribution \mathbf{p} . Let \mathbf{p}^* be the true input distribution, and $\mu_j = E_{\mathbf{p}^*}[\phi_j(h(\mathbf{X}))]$ for a chosen sequence of moment functions ϕ_j , $j = 1, 2, \dots$. Assume the following:

1. Let $\mathcal{N}(\mathbf{p}^*) \subset \mathcal{P}$ be a small neighborhood of \mathbf{p}^* . For every $\mathbf{p} \in \mathcal{N}(\mathbf{p}^*)$, $F_Y(\mathbf{p})$ is completely identified by a finite number of ϕ_j , i.e. $F_Y(\mathbf{p})$ is uniquely determined by

$$\int \phi_k(y) dF_Y(\mathbf{p})(y) = \mu_j, j = 1, \dots, l$$

among the set of distributions $\mathcal{P}_Y = \{F_Y(\mathbf{q}) : \mathbf{q} \in \mathcal{P}\}$, for some $l > 0$. This defines a map $\mathcal{G} : \mathcal{N}(\mathbf{p}^*) \rightarrow \mathcal{N}(\boldsymbol{\mu})$, where $\mathcal{N}(\boldsymbol{\mu})$ denotes a small neighborhood of $\boldsymbol{\mu} = (\mu_1, \dots, \mu_l) \in \mathbb{R}^l$. We assume moreover that \mathcal{G} is one-to-one.

2. \mathcal{G}^{-1} , the inverse of \mathcal{G} , is continuous.

3. The empirical moment $\hat{\mu}_j \rightarrow \mu_j$ a.s. as the output sample size $N \rightarrow \infty$, for $j = 1, \dots, l$.

Then $\mathcal{G}^{-1}(\hat{\boldsymbol{\mu}}) \rightarrow \mathbf{p}^*$ a.s. as $N \rightarrow \infty$, where $\hat{\boldsymbol{\mu}} = (\hat{\mu}_1, \dots, \hat{\mu}_l)$.

Proof of Lemma 1. As $N \rightarrow \infty$, by Assumption 3, we have $\hat{\mu}_j \rightarrow \mu_j$. Hence by Assumptions 1 and 2, we have $\mathcal{G}^{-1}(\hat{\boldsymbol{\mu}}) \rightarrow \mathbf{p}^*$. \square

Lemma 1 is of course a rather trivial result, stating that with enough moments and a one-to-one map from the input to the output distribution, one would be able to consistently calibrate the input model. The conditions in the lemma are generally very difficult to verify in practice, since the map \mathcal{G} is related to the simulation process that could be highly complex. The use of Lemma 1 is more about pointing out the limitation of moment-matching: Roughly speaking, we expect that, if the input support size is n , we need $l = n - 1$ moments to form a well-determined system of equations and subsequently recover the input model. Hence when n is big, the number of moments matched needs also be big. However, with finite output observation size, statistical errors prohibit matching too many moments, since in this case some of the empirical moments could be badly estimated (i.e., overfitting). Thus in practice one would only match a small number of moments. This in turn implies that one would have an under-determined system of equations that results in plenty of candidate input models. An additional criterion, namely the entropy, is therefore needed to pin down the choices.

3.2. Why Maximize Entropy?

The use of the ME criterion is strongly inspired from results in density estimation. Taking the simulation model aside, and consider estimating the density from a finite number of i.i.d. data. It is known that by increasing the number of matched moments suitably with the sample size, the ME density subject to empirical moment matching is a consistent density estimator and, moreover, elicits its explicitly obtainable convergence rate to the true density, in terms of KL divergence, according to the class of moment functions used (Barron and Sheu (1991)). The only difference between our setting and this classical setting is the layer of transformation induced by the simulation model. We note that one could potentially derive the precise conditions on the simulation model to generalize these classical results. However, such an attempt may not be meaningful because the obtained conditions are likely unverifiable for typically complex simulation maps. We thus choose to only justify the ME criterion on an intuitive level.

In addition, entropy has been widely used as a proxy for the expected gain or information in experimental design (e.g., Lindley (1956), Box and Hill (1967), Bernardo (1979)). The entropy here often refers to the output distribution, though input distribution has also been considered if the goal is to optimize inference at the input level (Chick and Ng (2002)). ME is also used in financial option pricing to infer the most “natural” risk-neutral measure nonparametrically (Avellaneda et al. (2001)). This paper, in some sense, follows these widely taken viewpoints of ME, under the presumption that we do not have any additional information about the input model.

4. Optimization Procedure

This section focuses on solving our main formulation (3). Our strategy consists of two parts: a reformulation of (3) into a sequence of optimization programs with deterministic convex constraints, which is inspired from the quadratic penalty method in nonlinear programming (Section 4.1), and a constrained SA procedure for finding local optima for each of these programs (Section 4.2).

4.1. Transforming into a Sequence of Stochastic Programs with Entropy Constraints

Note that the constraints in (3) are in general nonlinear because the i.i.d. input sequence means that the expectation $E_{\mathbf{p}}[\cdot]$ is a convolution of \mathbf{p} . In fact, $E_{\mathbf{p}}[\phi_j(h(\mathbf{X}))]$ is a high-dimensional polynomial in \mathbf{p} , and is in general non-convex. Moreover, this polynomial can involve a huge (or even infinite) number of terms and hence its evaluation requires simulation approximation.

To handle such non-convex stochastic constraints, consider a sequence of optimization programs parametrized by η

$$\begin{aligned} & \min && \sum_{j=1}^m (E_{\mathbf{p}}[\phi_j(h(\mathbf{X}))] - \hat{\mu}_j)^2 \\ & \text{subject to} && R(\mathbf{p}) \geq \eta \\ & && \mathbf{p} \in \mathcal{P} \end{aligned} \tag{4}$$

This sequence satisfies the following key properties:

THEOREM 1. *Suppose that $E_{\mathbf{p}}[\phi_j(h(\mathbf{X}))]$ is a continuous function in \mathbf{p} , for each $j = 1, \dots, m$. Let $W^*(\eta)$ denotes the optimal value of (4) indexed at η . We have*

1. $W^*(\eta)$ decreases as η decreases from $\log n$ to 0.
2. An optimal solution of (4), at any $\eta \in [0, \log n]$, exists. Denote this optimal solution as $\mathbf{p}^*(\eta)$.
3. If there exists an $\eta^* = \sup\{\eta \in [0, \log n] : W^*(\eta) = 0\}$, then $\mathbf{p}^*(\eta^*)$ is optimal for (3).
4. If there does not exist any $\eta \in [0, \log n]$ such that $W^*(\eta) = 0$, then (3) is infeasible.

Proof of Theorem 1 *Proof of 1.* First, it is elementary to check that $0 \leq R(\mathbf{p}) \leq \log n$ for any $\mathbf{p} \in \mathcal{P}$, with the lower bound achieved by a point mass on any support point in \mathcal{S} , and the upper bound achieved by a uniform distribution on \mathcal{S} . It is also routine to check that there exists \mathbf{p} that has $R(\mathbf{p}) = r$ for any $r \in [0, \log n]$. Thus program (4) is feasible for any $\eta \in [0, \log n]$. As η decreases, the feasible region in (4) enlarges, and so the optimal value of (4) decreases.

Proof of 2. Note that the feasible region of (4) is compact for any η . Since, by the assumption, the objective function is continuous in \mathbf{p} , $\mathbf{p}^*(\eta)$ exists by the Weierstrass Theorem.

Proof of 3. Suppose that an $\eta^* = \sup\{\eta \in [0, \log n] : W^*(\eta) = 0\}$ exists. Then by definition there must exist an η such that $W^*(\eta) = 0$, which implies, by part (b), that there is a feasible solution for (3) with $R(\mathbf{p}) \geq \eta$. Thus (3) is feasible. As the intersection of a simplex and the set of roots of continuous functions, the feasible region of (3) is compact, and also its objective function is continuous. Hence by the Weierstrass Theorem again, (3) possesses an optimal solution $\tilde{\mathbf{p}}^*$. We denote $\tilde{\eta}^* = R(\tilde{\mathbf{p}}^*)$ as the optimal value of (3).

Consider $\tilde{\eta}^* + \epsilon$ for $\epsilon > 0$. By the definition of $\tilde{\eta}^*$ there does not exist any feasible solution \mathbf{p} for (3) with $R(\mathbf{p}) \geq \tilde{\eta}^* + \epsilon$, and hence $W^*(\tilde{\eta}^* + \epsilon) > 0$.

On the other hand, consider $\tilde{\eta}^* - \epsilon$ for $\epsilon \geq 0$. We have $W^*(\tilde{\eta}^* - \epsilon) = 0$ by observing that $\tilde{\mathbf{p}}^*$ is a solution for (4) at $\tilde{\eta}^* - \epsilon$.

Therefore, $\tilde{\eta}^* = \eta^*$ and $W^*(\eta^*) = W^*(\tilde{\eta}^*) = 0$. This implies that $\mathbf{p}^*(\eta^*)$ is feasible and has objective value at least $\eta^* = \tilde{\eta}^*$ for (3). By the definition of $\tilde{\eta}^*$, $\mathbf{p}^*(\eta^*)$ is optimal and has objective value η^* for (3).

Proof of 4. If there does not exist $\eta \in [0, \log n]$ such that $W^*(\eta) = 0$, then there is no $\mathbf{p} \in \mathcal{P}$ that satisfies $E_{\mathbf{p}}[\phi_j(h(\mathbf{X}))] = \hat{\mu}_j$ for all $j = 1, \dots, m$, and (3) is infeasible. \square

In view of Theorem 1, our strategy is to solve (4) at different values of η and identify η^* . This strategy is a modification of the quadratic penalty method for solving deterministic nonlinear programs, where equality constraints are relaxed into the objective function with squaring (Bertsekas (1999)). Note that while the standard quadratic penalty method solves a sequence of optimizations with objectives consisting of both the primal objective and the relaxed squared constraints, our formulation (4) has chosen to put the primal objective as a constraint in the optimization

sequence. This arrangement is beneficial for a few reasons. First, it allows us to more easily locate a stochastic root $\eta^* \in [0, \log n]$ for which we can look for an optimal solution, whereas the standard quadratic penalty method requires the sequencing index to go to ∞ and does not offer an explicit guideline on when to stop searching. Second, our primal objective, as an entropy, has advantageous concave structure that allows running efficient MDSA (which we will discuss momentarily) when translated as constraint. Third, η has the interpretation as the entropy level of the distributions to be considered. For convenience we call our method the *stochastic quadratic penalty method*. In the Supplementary Materials, we provide some discussion on applying a variant of our approach that is closer to the conventional quadratic penalty method.

4.2. Constrained Stochastic Approximation for Solving the Optimization Sequence

Although formulation (4) is still non-convex, its constraints are convex and deterministic, which can be handled more easily using SA than in the original formulation (3). This section investigates the design and analysis of an MDSA algorithm for finding a local optimum of (4).

MDSA is the stochastization of the mirror descent (MD) method, an iterative procedure for solving deterministic convex programs (Nemirovski and Yudin (1983)). MD was first motivated for optimizations in general normed space, where the space of the solution (i.e. the primal) is different from the space of the gradient (i.e. the dual), and hence the usual gradient descent scheme does not make sense. MD uses the insight of first mapping the solution to the dual space (using a map associated with a strongly convex function called the distance-generating function), moving the solution along the gradient direction in the dual space, and mapping it back to the primal (via the so-called prox-mapping). In the Euclidean space, it is not necessary to use MD since the primal and the dual space are the same. But, by choosing a suitable primal-dual mapping, MD can provide advantages in the convergence speed in terms of less dependence on the problem dimension. Moreover, the resulting optimization subroutine required in the iteration can be very efficient.

To describe the algorithm, MD finds the next iterate via optimizing the objective function linearized at the current solution, together with a penalty on the distance of movement of the next iterate. When the objective function is only accessible via simulation, the linearized objective function, or the gradient, at each iteration can only be estimated, in which case the procedure becomes MDSA (Nemirovski et al. (2009)). More precisely, given a current iterate \mathbf{p}^k , MDSA solves

$$\begin{aligned} \min \quad & \gamma^k \hat{\boldsymbol{\psi}}^k{}'(\mathbf{p} - \mathbf{p}^k) + V(\mathbf{p}^k, \mathbf{p}) \\ \text{subject to } & R(\mathbf{p}) \geq \eta \\ & \mathbf{p} \in \mathcal{P} \end{aligned} \tag{5}$$

where $\hat{\boldsymbol{\psi}}^k$ carries the gradient information at \mathbf{p}^k , $V(\cdot, \cdot)$ is some distance measure known as the prox-function (Nemirovski et al. (2009)), and γ^k is the step size at iteration k . (5) thus minimizes

over the feasible set of \mathbf{p} with an objective function linearized at \mathbf{p}^k , penalized by the distance $(1/\gamma^k)V(\mathbf{p}^k, \cdot)$. To implement this scheme, we need to investigate: 1) how to obtain $\hat{\psi}^k$, 2) the complexity of the program (5) with a choice of V , and 3) the convergence property of the procedure in relation to γ^k . The next three subsections present these investigations respectively.

Even though MDSA can provably converge to a local optimum, the non-convexity of (4) means that there is no guarantee of finding a global one. However, we can obtain some evidence of global convergence by scrutinizing the monotonic pattern of $W^*(\eta)$ as predicted by Theorem 1. We will revisit this discussion in our numerical experiments in Section 5.

4.2.1. Gradient Estimation. We denote $W(\mathbf{p})$ as the objective function in (4). Though $W(\mathbf{p})$ is a function on the Euclidean space and in principle can be differentiated in a standard manner, naive differentiation of $W(\mathbf{p})$ with respect to \mathbf{p} in general does not lead to any simulable form. This is because an arbitrary perturbation of \mathbf{p} can shoot outside the probability simplex, and the resulting gradient will be a high-dimensional polynomial in \mathbf{p} that has no probabilistic interpretation. To get around this issue, we use the idea of the Gateaux derivative defined on a functional of probability distribution (Serfling (2009)). This consists of restricting the perturbations of \mathbf{p} within the probability simplex as represented by the mixtures $(1 - \epsilon)\mathbf{p} + \epsilon\mathbf{1}_i$, where $\mathbf{1}_i$ is a point mass at the support point z_i and $0 \leq \epsilon \leq 1$ is a mixture parameter. The idea is to differentiate $W((1 - \epsilon)\mathbf{p} + \epsilon\mathbf{1}_i)$ at $\epsilon = 0$. The resulting quantity, which we call $\psi_i(\mathbf{p})$, satisfies the following:

PROPOSITION 1. *We have:*

1. *Suppose W is differentiable in \mathcal{P} , then*

$$\nabla W(\mathbf{p})'(\mathbf{q} - \mathbf{p}) = \psi(\mathbf{p})'(\mathbf{q} - \mathbf{p}) \quad (6)$$

for any $\mathbf{q} \in \mathcal{P}$, where $\psi(\mathbf{p}) = (\psi_1(\mathbf{p}), \dots, \psi_n(\mathbf{p}))'$ and

$$\psi_i(\mathbf{p}) = \left. \frac{d}{d\epsilon} W((1 - \epsilon)\mathbf{p} + \epsilon\mathbf{1}_i) \right|_{\epsilon=0^+}$$

2. *Assume, for all $j = 1, \dots, m$, $E_{\mathbf{p}}[|\phi_j(h(\mathbf{X}))|^{l+\theta}] < \infty$ for an integer $l \geq 1$ and some small $\theta > 0$, and τ satisfies $E_{\mathbf{p}}[e^{\beta\tau}] < \infty$ for some small $\beta > 0$. Also assume $\mathbf{p} = (p_1, \dots, p_n)$ where each $p_i > 0$. θ, β might depend on \mathbf{p} . Then $\psi_i(\mathbf{p})$ is finite for all i and is equal to*

$$\psi_i(\mathbf{p}) = 2 \sum_{j=1}^m (E_{\mathbf{p}}[\phi_j(h(\mathbf{X}))] - \hat{\mu}_j) E_{\mathbf{p}}[\phi_j(h(\mathbf{X})) S_i(\mathbf{X}; \mathbf{p})] \quad (7)$$

$$= 2E_{\mathbf{p}} \left[\sum_{j=1}^m (\phi_j(h(\mathbf{X})) - \hat{\mu}_j) \phi_j(h(\tilde{\mathbf{X}})) S_i(\tilde{\mathbf{X}}; \mathbf{p}) \right] \quad (8)$$

where

$$S_i(\mathbf{x}; \mathbf{p}) = \sum_{t=1}^{\tau} \frac{I_i(x_t)}{p_i} - \tau$$

for $\mathbf{x} = (x_1, \dots, x_\tau)$. Here $I_i(x) = 1$ if $x = z_i$ and 0 otherwise. \mathbf{X} and $\tilde{\mathbf{X}}$ are two independent copies of the i.i.d. input process generated under \mathbf{p} . Moreover, for all i, j and $1 \leq s \leq l$, the moments

$$E_{\mathbf{p}}[(\phi_j(h(\mathbf{X})))^s], E_{\mathbf{p}}[(\phi_j(h(\mathbf{X}))S_i(\mathbf{X}; \mathbf{p}))^s]$$

are continuous in \mathbf{p} .

(6) guarantees that the Gateaux derivative $\psi(\mathbf{p})$ behaves the same as a standard gradient $\nabla W(\mathbf{p})$ when applying to any directions within the probability simplex \mathcal{P} , which are the only directions we consider in MDSA. Importantly, $\psi(\mathbf{p})$ is simulable by using (8). Through (7) and (8), each component $\psi_i(\mathbf{p})$ can now be expressed via the function $S_i(\cdot; \mathbf{p})$ that plays the role of the score function as in the conventional likelihood ratio method (also known as the score function method) (Glynn (1990), Reiman and Weiss (1989), Rubinstein (1989), L'Ecuyer (1990)) in parametric sensitivity analysis. Thus Proposition 1 can be viewed as a nonparametric version of the likelihood ratio method. Ghosh and Lam (2015a) and Ghosh and Lam (2015b) have also studied such type of methods for finding the gradients of expectation-type performance measures. Proposition 1 Part 2 generalizes their results to sums of squared expectations and random time horizons, relaxes their boundedness condition on the performance function, and identifies the conditions to guarantee finite moments of the estimator and their continuity with respect to the underlying probability measure.

Proof of Proposition 1. Proof of 1. We have

$$\left. \frac{d}{d\epsilon} W((1-\epsilon)\mathbf{p} + \epsilon\mathbf{1}_k) \right|_{\epsilon=0^+} = \nabla W(\mathbf{p})'(\mathbf{1}_k - \mathbf{p}) = \frac{\partial}{\partial p_i} W(\mathbf{p}) - \nabla W_j(\mathbf{p})'\mathbf{p}$$

and so

$$\psi(\mathbf{p}) = \nabla W(\mathbf{p}) - (\nabla W(\mathbf{p})'\mathbf{p})\mathbf{1}$$

where $\mathbf{1}$ is a vector of 1's. Therefore,

$$\psi(\mathbf{p})'(\mathbf{q} - \mathbf{p}) = \nabla W(\mathbf{p})'(\mathbf{q} - \mathbf{p}) - (\nabla W(\mathbf{p})'\mathbf{p})(\mathbf{1} - \mathbf{1}) = \nabla W(\mathbf{p})'(\mathbf{q} - \mathbf{p})$$

which concludes (6).

Proof of 2. For convenience, let $u_j(\mathbf{p}) = E_{\mathbf{p}}[\phi_j(h(\mathbf{X}))]$. Fix any i . Using the chain rule, we have

$$\left. \frac{d}{d\epsilon} W((1-\epsilon)\mathbf{p} + \epsilon\mathbf{1}_i) \right|_{\epsilon=0} = 2 \sum_{j=1}^m (u_j(\mathbf{p}) - \hat{\mu}_j) \frac{d}{d\epsilon} u_j((1-\epsilon)\mathbf{p} + \epsilon\mathbf{1}_i) \Big|_{\epsilon=0}$$

We argue that the derivative $\frac{d}{d\epsilon}u_j((1-\epsilon)\mathbf{p} + \epsilon\mathbf{1}_i)|_{\epsilon=0}$ exists and is equal to $E_{\mathbf{p}}[\phi_j(h(\mathbf{X}))S_i(\mathbf{X};\mathbf{p})]$. Note that by our assumption $p_i \geq \delta$ for all i for some $\delta > 0$. Thus there exists a small enough ν such that $(1-\epsilon)\mathbf{p} + \epsilon\mathbf{1}_i$ lies in \mathcal{P} for all $\epsilon \in (-\nu, \nu)$. For any $x \in \mathcal{S}$, denote $p(x)$ as the probability mass on x . We rewrite the expectation as one with respect to \mathbf{p} by inserting the likelihood ratio

$$u_j((1-\epsilon)\mathbf{p} + \epsilon\mathbf{1}_i) = E_{\mathbf{p}} \left[\phi_j(h(\mathbf{X})) \prod_{t=1}^{\tau} \frac{(1-\epsilon)p(x_t) + \epsilon\mathbf{1}_i(x_t)}{p(x_t)} \right].$$

Hence

$$\begin{aligned} & \frac{u_j((1-\epsilon)\mathbf{p} + \epsilon\mathbf{1}_i) - u_j(\mathbf{p})}{\epsilon} \\ &= E_{\mathbf{p}} \left[\phi_j(h(\mathbf{X})) \prod_{t=1}^{\tau} \frac{(1-\epsilon')p(x_t) + \epsilon'\mathbf{1}_i(x_t)}{p(x_t)} \sum_{t=1}^{\tau} \frac{\mathbf{1}_i(x_t) - p(x_t)}{(1-\epsilon')p(x_t) + \epsilon'\mathbf{1}_i(x_t)} \right] \end{aligned} \quad (9)$$

where we use mean value theorem for the product, and $\epsilon' \in [0, \epsilon]$. The product and sum in (9) can be bounded as

$$\begin{aligned} & \prod_{t=1}^{\tau} \frac{(1-\epsilon')p(x_t) + \epsilon'\mathbf{1}_i(x_t)}{p(x_t)} = \prod_{t=1}^{\tau} \frac{p(x_t) + \epsilon'(\mathbf{1}_i(x_t) - p(x_t))}{p(x_t)} \leq \left(1 + \frac{\nu}{\delta}\right)^{\tau} \\ & \left| \sum_{t=1}^{\tau} \frac{\mathbf{1}_i(x_t) - p(x_t)}{(1-\epsilon')p(x_t) + \epsilon'\mathbf{1}_i(x_t)} \right| \leq \sum_{t=1}^{\tau} \frac{1}{|p(x_t) + \epsilon'(\mathbf{1}_i(x_t) - p(x_t))|} \leq \frac{\tau}{\delta - \nu} \end{aligned}$$

Also by Holder's inequality

$$E_{\mathbf{p}} \left[|\phi_j(h(\mathbf{X}))| \left(1 + \frac{\nu}{\delta}\right)^{\tau} \frac{\tau}{\delta - \nu} \right] \leq \left(E_{\mathbf{p}}[|\phi_j(h(\mathbf{X}))|^{k_1}]\right)^{\frac{1}{k_1}} \left(E_{\mathbf{p}} \left[\left(1 + \frac{\nu}{\delta}\right)^{k_2\tau} \left(\frac{\tau}{\delta - \nu}\right)^{k_2} \right]\right)^{\frac{1}{k_2}}, \quad (10)$$

where $1 < k_1 < l + \theta, k_2 > 1, \frac{1}{k_1} + \frac{1}{k_2} = 1$. The right side of (10) is finite if ν is chosen such that $(1 + \frac{\nu}{\delta})^{k_2} < e^{\beta}, \nu < \delta$. Therefore, by the dominated convergence theorem

$$\lim_{\epsilon \rightarrow 0} \frac{u_j((1-\epsilon)\mathbf{p} + \epsilon\mathbf{1}_i) - u_j(\mathbf{p})}{\epsilon} = E_{\mathbf{p}} \left[\phi_j(h(\mathbf{X})) \sum_{t=1}^{\tau} \frac{\mathbf{1}_i(x_t) - p(x_t)}{p(x_t)} \right].$$

This proves the first half of the statement. To prove the second half, fix a \mathbf{p} . When examining the continuity at \mathbf{p} , we always change the underlying probability measure to \mathbf{p}

$$E_{\mathbf{q}}[(\phi_j(h(\mathbf{X}))S_i(\mathbf{X};\mathbf{q}))^s] = E_{\mathbf{p}} \left[(\phi_j(h(\mathbf{X})))^s \left(\sum_{t=1}^{\tau} \frac{\mathbf{1}_i(x_t) - q(x_t)}{q(x_t)} \right)^s \prod_{t=1}^{\tau} \frac{q(x_t)}{p(x_t)} \right]. \quad (11)$$

Similarly, we bound the product and sum in (11) as

$$\begin{aligned} & \prod_{t=1}^{\tau} \frac{q(x_t)}{p(x_t)} \leq \left(\max_i \frac{q_i}{p_i} \right)^{\tau} \\ & \left| \sum_{t=1}^{\tau} \frac{\mathbf{1}_i(x_t) - q(x_t)}{q(x_t)} \right| \leq \frac{\tau}{\min_i q_i}. \end{aligned}$$

Note that when \mathbf{q} is in a small enough neighborhood of \mathbf{p} , $\max_i \frac{q_i}{p_i}$ can be arbitrarily close to one, say less than $1 + \epsilon_r$, and $\min_i q_i$ is uniformly bounded below by some $\delta_r > 0$. So by Holder's inequality

$$E_{\mathbf{p}} \left[|\phi_j(h(\mathbf{X}))|^s \frac{\tau^s}{\delta_r^s} (1 + \epsilon_r)^\tau \right] \leq \left(E_{\mathbf{p}} [|\phi_j(h(\mathbf{X}))|^{k_1 s}] \right)^{\frac{1}{k_1}} \left(E_{\mathbf{p}} \left[\frac{\tau^{k_2 s}}{\delta_r^{k_2 s}} (1 + \epsilon_r)^{k_2 \tau} \right] \right)^{\frac{1}{k_2}}.$$

where $1 < k_1 < \frac{l+\theta}{s}$, $k_2 > 1$, $\frac{1}{k_1} + \frac{1}{k_2} = 1$. As in the proof of differentiability, the last expectation is finite when $(1 + \epsilon_r)^{k_2} < e^\beta$. By dominated convergence theorem and (11)

$$\lim_{\mathbf{q} \rightarrow \mathbf{p}} E_{\mathbf{q}} [(\phi_j(h(\mathbf{X})) S_i(\mathbf{X}; \mathbf{q}))^s] = E_{\mathbf{p}} [(\phi_j(h(\mathbf{X})) S_i(\mathbf{X}; \mathbf{p}))^s].$$

The proof of continuity of $E_{\mathbf{p}} [(\phi_j(h(\mathbf{X})))^s]$ is similar. \square

In view of Proposition 1, we have the following gradient estimation scheme:

COROLLARY 1. *Under the assumptions in Proposition 1 Part 2, an unbiased estimator for $\psi(\mathbf{p})$ is given by $\widehat{\psi}(\mathbf{p}) = (\widehat{\psi}_1(\mathbf{p}), \dots, \widehat{\psi}_n(\mathbf{p}))'$, where*

$$\widehat{\psi}_i(\mathbf{p}) = 2 \sum_{j=1}^m \frac{1}{M_1} \sum_{r=1}^{M_1} (\phi_j(h(\mathbf{X}^{(r)})) - \hat{\mu}_j) \frac{1}{M_2} \sum_{r=1}^{M_2} \phi_j(h(\tilde{\mathbf{X}}^{(r)})) S_i(\tilde{\mathbf{X}}^{(r)}; \mathbf{p}) \quad (12)$$

and $\mathbf{X}^{(r)}$'s and $\tilde{\mathbf{X}}^{(r)}$'s are M_1 and M_2 independent copies of the i.i.d. input process generated under \mathbf{p} and are used simultaneously in all $\widehat{\psi}_i(\mathbf{p})$.

4.2.2. Solving Stepwise Subprograms in MDSA. We discuss the choice of the prox-function V and how to solve program (5). Following Nemirovski et al. (2009), we take V as the Bregman divergence

$$V(\mathbf{p}, \mathbf{q}) = \omega(\mathbf{q}) - \omega(\mathbf{p}) - \nabla \omega(\mathbf{p})'(\mathbf{q} - \mathbf{p}) \quad (13)$$

where $\omega(\cdot)$ is called the distance-generating function and is strongly convex, i.e.

$$\omega(\mathbf{q}) - \omega(\mathbf{p}) \geq \nabla \omega(\mathbf{p})'(\mathbf{q} - \mathbf{p}) + \frac{\alpha}{2} \|\mathbf{q} - \mathbf{p}\|^2 \quad (14)$$

for all \mathbf{p}, \mathbf{q} in the feasible region, $\|\cdot\|$ is some norm, and $\alpha > 0$. One choice of ω that is especially appropriate for our feasible space, the probability simplex, is the negative entropy $\omega(\mathbf{p}) = \sum_{i=1}^m p_i \log p_i$. In this case $\alpha = 1$ and $\|\cdot\|$ is taken as the L_1 -norm. The function $V(\mathbf{p}, \mathbf{q})$ derived from (13) is the KL divergence given by

$$V(\mathbf{p}, \mathbf{q}) = \sum_{i=1}^n q_i \log \frac{q_i}{p_i} \quad (15)$$

We shall use this choice in our procedure. The iteration (5) using (15), but without the constraint $R(\mathbf{p}) \geq \eta$ and the stochasticity, is also called the entropic descent algorithm (Beck and Teboulle (2003)).

Consider a generic formulation of (5) written as

$$\begin{aligned} & \min \quad \boldsymbol{\xi}'(\mathbf{q} - \mathbf{p}) + V(\mathbf{p}, \mathbf{q}) \\ & \text{subject to } R(\mathbf{q}) \geq \eta \\ & \quad \mathbf{q} \in \mathcal{P} \end{aligned} \quad (16)$$

where $V(\mathbf{p}, \mathbf{q})$ is defined by (15), for some given $\mathbf{p} = (p_1, \dots, p_n)$ and $\boldsymbol{\xi} = (\xi_1, \dots, \xi_n)$. The solution of (16) is in the following form:

PROPOSITION 2. Denote

$$\kappa(\mathbf{p}, \boldsymbol{\xi}, \alpha) = \frac{\sum_{i=1}^n p_i^{\frac{1}{1+\alpha}} e^{-\frac{\xi_i}{1+\alpha}} \left(-\frac{1}{1+\alpha} \log p_i + \frac{\xi_i}{1+\alpha} \right)}{\sum_{i=1}^n p_i^{\frac{1}{1+\alpha}} e^{-\frac{\xi_i}{1+\alpha}}} + \log \sum_{i=1}^n p_i^{\frac{1}{1+\alpha}} e^{-\frac{\xi_i}{1+\alpha}} \quad (17)$$

Suppose $\kappa(\mathbf{p}, \boldsymbol{\xi}, 0) \geq \eta$, then the optimal solution of (16) is given by $\mathbf{q}^* = (q_1^*, \dots, q_n^*)$, where

$$q_i^* = \frac{p_i e^{-\xi_i}}{\sum_{l=1}^n p_l e^{-\xi_l}} \quad (18)$$

Otherwise, if $\kappa(\mathbf{p}, \boldsymbol{\xi}, 0) < \eta$, then

$$q_i^* = \frac{p_i^{\frac{1}{1+\alpha^*}} e^{-\frac{\xi_i}{1+\alpha^*}}}{\sum_{l=1}^n p_l^{\frac{1}{1+\alpha^*}} e^{-\frac{\xi_l}{1+\alpha^*}}} \quad (19)$$

where α^* is a positive root (potentially ∞) of the equation $\kappa(\mathbf{p}, \boldsymbol{\xi}, \alpha) = \eta$.

Proof of Proposition 2 Consider the Lagrangian

$$\begin{aligned} & \boldsymbol{\xi}'(\mathbf{q} - \mathbf{p}) + V(\mathbf{p}, \mathbf{q}) - \alpha(R(\mathbf{q}) - \eta) + \lambda \left(\sum_{i=1}^n q_i - 1 \right) \\ & = \boldsymbol{\xi}'(\mathbf{q} - \mathbf{p}) + \sum_{i=1}^n q_i \log \frac{q_i}{p_i} + \alpha \sum_{i=1}^n q_i \log q_i + \alpha \eta + \lambda \left(\sum_{i=1}^n q_i - 1 \right) \end{aligned} \quad (20)$$

for $\alpha \geq 0$ and $\lambda \in \mathbb{R}$, by relaxing the constraints $R(\mathbf{q}) \geq \eta$ and $\sum_{i=1}^n q_i = 1$.

Suppose $\kappa(\mathbf{p}, \boldsymbol{\xi}, 0) \geq \eta$. Consider $\alpha = 0$. Differentiating (20) with respect to \mathbf{q} gives

$$\xi_i + \log \frac{q_i}{p_i} + 1 + \lambda$$

Setting it to zero, we have

$$q_i \propto p_i e^{-\xi_i}$$

Using the constraint $\sum_{i=1}^n q_i = 1$, we get \mathbf{q}^* defined in (18). Note that $R(\mathbf{q}^*) = \kappa(\mathbf{p}, \boldsymbol{\xi}, 0) \geq \eta$. It is then straightforward to verify that \mathbf{q}^* satisfies the KKT condition.

Suppose now $\kappa(\mathbf{p}, \boldsymbol{\xi}, 0) < \eta$. Consider $\alpha > 0$. Differentiating (20) gives

$$\xi_i + \log \frac{q_i}{p_i} + 1 + \alpha \log q_i + \alpha + \lambda$$

and setting to zero gives

$$q_i \propto p_i^{\frac{1}{1+\alpha}} e^{-\frac{\xi_i}{1+\alpha}}$$

Using the constraint $\sum_{i=1}^n q_i = 1$, and setting $R(\mathbf{q}^*) = \eta$, we get \mathbf{q}^* defined in (19) since now $R(\mathbf{q}^*) = \kappa(\mathbf{p}, \boldsymbol{\xi}, \alpha)$. Moreover, for $\eta < \log n$, the equation $\kappa(\mathbf{p}, \boldsymbol{\xi}, \alpha) = \eta$ must have a positive root because $\kappa(\mathbf{p}, \boldsymbol{\xi}, 0) < \eta$ in the current case, and $\kappa(\mathbf{p}, \boldsymbol{\xi}, \alpha) \rightarrow \log n > \eta$. It is then straightforward to verify that \mathbf{q}^* satisfies the KKT condition. Finally, when $\eta = \log n$, we have $\alpha^* = \infty$ and \mathbf{q}^* is a uniform distribution, which is consistent with (19). \square

Proposition 2 entails that the subprogram in each iteration can be solved as a one-dimensional root-finding problem, which can be implemented efficiently by, e.g., a bisection search.

4.2.3. Convergence Analysis. The MDSA algorithm is depicted in Algorithm 1. Steps 1 and 2 there come from the discussion in Sections 4.2.1 and 4.2.2 respectively. Step 3 is a technical step in securing theoretical convergence (as will be shown momentarily). We point out that Steps 2 and 3 combined are in effect solving (5) with \mathcal{P} replaced by $\mathcal{P}(\epsilon)$, a restricted version of the original (5). Step 2 first solves (5) under \mathcal{P} . If its optimal solution lies in $\mathcal{P}(\epsilon)$, then this is immediately an optimal solution for the restricted problem. Otherwise, Step 3 is carried out to find the optimal solution for the restricted problem directly.

The reason why we consider such a restricted problem is to guarantee that \mathbf{p}^k does not have any components that are too small. In turn, this is because the form of the gradient estimator $\hat{\boldsymbol{\psi}}^k$ contains p_i^k at the denominator, and a small p_i^k can blow up its variance. By restricting p_i^k to be at least ϵ , the variance of $\hat{\boldsymbol{\psi}}^k$ is bounded, and convergence of Algorithm 1 can be shown in situations where the optimal solution for (4) is indeed in $\mathcal{P}(\epsilon)$.

We mention that Step 3 is really a technicality for theoretical correctness and does not seem to have practical implications. All the experiments we perform (in Section 5) do not run into the problem of vanishing p_i^k . For this reason we do not attempt to find analytical solution for the optimization (5) under $\mathcal{P}(\epsilon)$ but rather just impose it as a general convex optimization problem.

THEOREM 2. *Suppose the assumptions in Proposition 1 Part 2 hold with $l = 2$. Assume there exists a unique optimal solution $\mathbf{p}^* \in \mathcal{P}(\epsilon)$ for (4) such that $\boldsymbol{\psi}(\mathbf{p})'(\mathbf{p} - \mathbf{p}^*) = 0$ if and only if $\mathbf{p} = \mathbf{p}^*$. Choose the step size sequence $\{\gamma^k\}$ such that*

$$\sum_{k=1}^{\infty} \gamma^k = \infty, \quad \sum_{k=1}^{\infty} (\gamma^k)^2 < \infty$$

Then \mathbf{p}^k generated in Algorithm 1 converges to \mathbf{p}^ a.s..*

The condition $\boldsymbol{\psi}(\mathbf{p})'(\mathbf{p} - \mathbf{p}^*) = 0$ is a generalization of the first order local optimality condition $\boldsymbol{\psi}(\mathbf{p}) = \mathbf{0}$ in unconstrained optimization. It is in line with the standard condition $\boldsymbol{\psi}(\mathbf{p})'(\mathbf{p} - \mathbf{p}^*) > 0$ for all $\mathbf{p} \neq \mathbf{p}^*$ used in the SA literature (e.g., Benveniste et al. (2012), Broadie et al. (2011)).

Algorithm 1 MDSA for solving (4)

Input: A small parameter $\epsilon > 0$, initial solution $\mathbf{p}^1 \in \mathcal{P}(\epsilon) = \{\mathbf{p} : \sum_{i=1}^n p_i = 1, p_i \geq \epsilon \text{ for } i = 1, \dots, n\}$, a step size sequence γ^k , and sample sizes M_1 and M_2 .

Iteration: For $k = 1, 2, \dots$, do the following: Given \mathbf{p}^k ,

1. Estimate $\hat{\boldsymbol{\psi}}^k = (\hat{\psi}_1^k, \dots, \hat{\psi}_n^k)$ with

$$\hat{\psi}_i^k = 2 \sum_{j=1}^m \frac{1}{M_1} \sum_{r=1}^{M_1} (\phi_j(h(\mathbf{X}^{(r)})) - \mu_j) \frac{1}{M_2} \sum_{r=1}^{M_2} \phi_j(h(\tilde{\mathbf{X}}^{(r)})) S_i(\tilde{\mathbf{X}}^{(r)}; \mathbf{p}^k)$$

where $\mathbf{X}^{(r)}$ and $\tilde{\mathbf{X}}^{(r)}$ are M_1 and M_2 independent copies of the input process generated under i.i.d. replications of \mathbf{p}^k , which are used simultaneously for all components of $\hat{\boldsymbol{\psi}}^k$.

2. If $\kappa(\mathbf{p}^k, \gamma^k \hat{\boldsymbol{\psi}}^k, 0) \geq \eta$ where κ is defined in (17), then output $\mathbf{p}^{k+1} = (p_1^{k+1}, \dots, p_n^{k+1})$, where

$$p_i^{k+1} = \frac{p_i^k e^{-\gamma^k \hat{\psi}_i^k}}{\sum_{l=1}^n p_l^k e^{-\gamma^k \hat{\psi}_l^k}}$$

Otherwise, if $\kappa(\mathbf{p}^k, \gamma^k \hat{\boldsymbol{\psi}}^k, 0) < \eta$, then

$$p_i^{k+1} = \frac{p_i^k \frac{1}{1+\alpha^k} e^{-\frac{\gamma^k \hat{\psi}_i^k}{1+\alpha^k}}}{\sum_{l=1}^n p_l^k \frac{1}{1+\alpha^k} e^{-\frac{\gamma^k \hat{\psi}_l^k}{1+\alpha^k}}}$$

where α^k is a positive root (potentially ∞) of the equation $\kappa(\mathbf{p}^k, \gamma^k \hat{\boldsymbol{\psi}}^k, \alpha) = \eta$.

3. If $p_i^{k+1} < \epsilon$ for some i , then solve the convex optimization (5) but with \mathcal{P} replaced by the set $\mathcal{P}(\epsilon)$. Output its solution as \mathbf{p}^{k+1} .

The proof of Theorem 2 follows the framework in Blum (1954), which considers SA on unconstrained problems.

Proof of Theorem 2 We borrow from Lemma 2.1 in Nemirovski et al. (2009) (restated in Lemma 2 in the Appendix) the inequality

$$V(\mathbf{p}^{k+1}, \mathbf{p}^*) - V(\mathbf{p}^k, \mathbf{p}^*) \leq \gamma^k \hat{\boldsymbol{\psi}}^k{}' (\mathbf{p}^* - \mathbf{p}^k) + \frac{\gamma^{k2} \|\hat{\boldsymbol{\psi}}^k\|_\infty^2}{2} \quad (21)$$

which holds as long as \mathbf{p}^{k+1} is the prox-mapping on \mathbf{p}^k obtained by solving (5). The norm $\|\cdot\|_\infty$ is the supremum norm, the dual of the L_1 -norm that is used in the strong convexity property (14) of $\omega(\mathbf{p}) = \sum_{i=1}^n p_i \log p_i$, with $\alpha = 1$.

Let \mathcal{F}^k be the filtration generated by $\{\mathbf{p}^1, \dots, \mathbf{p}^k\}$. Taking conditional expectation on (21), we have

$$E[V(\mathbf{p}^{k+1}, \mathbf{p}^*) - V(\mathbf{p}^k, \mathbf{p}^*) | \mathcal{F}^k] \leq \gamma^k \boldsymbol{\psi}(\mathbf{p}^k)' (\mathbf{p}^* - \mathbf{p}^k) + \frac{\gamma^{k2} E[\|\hat{\boldsymbol{\psi}}^k\|_\infty^2 | \mathcal{F}^k]}{2} \quad (22)$$

By the design of Algorithm 1, all components of \mathbf{p}^k is bounded from below by ϵ . To continue, we need a uniform bound of $E[\|\hat{\boldsymbol{\psi}}^k\|_\infty^2 | \mathcal{F}^k]$. From (12), Jensen's inequality and independence it follows

$$\begin{aligned} E[\|\hat{\boldsymbol{\psi}}^k\|_\infty^2 | \mathcal{F}^k] &\leq E\left[\sum_{i=1}^n \left(\hat{\psi}_i^k\right)^2 \middle| \mathcal{F}^k\right] \\ &\leq 4m \sum_{i=1}^n \sum_{j=1}^m E_{\mathbf{p}^k}[(\phi_j(h(\mathbf{X})) - \mu_j)^2] E_{\mathbf{p}^k}[(\phi_j(h(\mathbf{X})) S_i(\mathbf{X}; \mathbf{p}^k))^2]. \end{aligned}$$

By Proposition 1 Part 2, each summand is continuous in \mathbf{p}^k . Since \mathbf{p}^k is restricted to $\mathcal{P}(\epsilon)$, by compactness each summand is uniformly bounded by some $C_{i,j}$. Therefore $E[\|\hat{\boldsymbol{\psi}}^k\|_\infty^2 | \mathcal{F}^k] \leq C$ uniformly holds for some $C > 0$. Moreover, note that $\boldsymbol{\psi}(\mathbf{p}^k)'(\mathbf{p}^* - \mathbf{p}^k) \leq 0$ by the optimality of \mathbf{p}^* . Therefore, taking expectation and summing up on both sides of (22), we have

$$\begin{aligned} \sum_{k=1}^{\infty} E[E[V(\mathbf{p}^{k+1}, \mathbf{p}^*) - V(\mathbf{p}^k, \mathbf{p}^*) | \mathcal{F}^k]^+] &\leq \sum_{k=1}^{\infty} \gamma^k E[\boldsymbol{\psi}(\mathbf{p}^k)'(\mathbf{p}^* - \mathbf{p}^k)] + \sum_{k=1}^{\infty} \frac{\gamma^{k^2} E\|\hat{\boldsymbol{\psi}}^k\|_\infty^2}{2} \\ &\leq \frac{C}{2} \sum_{k=1}^{\infty} \gamma^{k^2} \\ &< \infty \end{aligned}$$

By martingale convergence theorem (Corollary in Section 3 in Blum (1954), restated in Theorem 3 in the Appendix), we have $V(\mathbf{p}^k, \mathbf{p}^*)$ converges a.s. to some integrable random variable V_∞ .

Now, taking expectation and summing up on both sides of (21), we have

$$\sum_{k=1}^{\infty} \gamma^k E[\boldsymbol{\psi}(\mathbf{p}^k)'(\mathbf{p}^k - \mathbf{p}^*)] \leq V(\mathbf{p}^1, \mathbf{p}^*) - EV_\infty + \sum_{k=1}^{\infty} \frac{\gamma^{k^2} E\|\hat{\boldsymbol{\psi}}^k\|_\infty^2}{2} < \infty$$

Since $\sum_{k=1}^{\infty} \gamma^k = \infty$ and $E[\boldsymbol{\psi}(\mathbf{p}^k)'(\mathbf{p}^k - \mathbf{p}^*)] \geq 0$, for all k , there must exist a subsequence k_i such that $E[\boldsymbol{\psi}(\mathbf{p}^{k_i})'(\mathbf{p}^{k_i} - \mathbf{p}^*)] \rightarrow 0$ a.s.. This implies that $\boldsymbol{\psi}(\mathbf{p}^{k_i})'(\mathbf{p}^{k_i} - \mathbf{p}^*) \xrightarrow{P} 0$, which in turn implies the existence of a further subsequence l_i such that $\boldsymbol{\psi}(\mathbf{p}^{l_i})'(\mathbf{p}^{l_i} - \mathbf{p}^*) \rightarrow 0$ a.s.. By Proposition 1 Part 2, we have $\boldsymbol{\psi}(\mathbf{p})$ continuous in \mathbf{p} . By the assumption that $\boldsymbol{\psi}(\mathbf{p})'(\mathbf{p} - \mathbf{p}^*) = 0$ only if $\mathbf{p} = \mathbf{p}^*$, and that $\boldsymbol{\psi}(\mathbf{p})'(\mathbf{p} - \mathbf{p}^*)$ is continuous in \mathbf{p} , we have $\mathbf{p}^{l_i} \rightarrow \mathbf{p}^*$ a.s.. Hence $V(\mathbf{p}^{l_i}, \mathbf{p}^*) \rightarrow 0$ a.s.. Since we have proved above that $V(\mathbf{p}^k, \mathbf{p}^*)$ converges a.s., this limit must be 0. Therefore, by Pinsker's inequality, we have $\mathbf{p}^k \rightarrow \mathbf{p}^*$ in total variation a.s.. This concludes the theorem. \square

5. Numerical Results

We provide numerical illustration of our method. We focus on a stylized M/G/1 queue, where we assume known i.i.d. exponential interarrival time distribution. Our goal is to calibrate the unknown i.i.d. service time distribution. In each experiment, we generate N i.i.d. realizations of $h(\mathbf{X})$ under a “true” service time distribution, for a specified function h . We take these N realizations as our

output data, and apply our formulation (3) and the stochastic quadratic penalty method in Section 4 to infer the service time distribution.

To apply the stochastic quadratic penalty method, we need to solve (4) for a grid of η such that η^* can be approximately located. We set the maximum search value of η to be the largest possible entropy $\log n$, and the minimum search value such that it is clearly far below the threshold η^* . This requires empirically keeping track of the approximate value of $W^*(\eta)$, the optimal value in (4), as we decrement η from $\log n$.

Denote by A_t the interarrival time between the t -th and $(t+1)$ -st customers, and by X_t the service time of the t -th customer. In all our examples shown below, we let the true service time be a discrete random variable over the set $\{1/n, 2/n, \dots, (n-1)/n, 1\}$, i.e. $z_i = i/n$, and $\mathbf{p} = (p_1, \dots, p_n)$ is obtained by

$$p_i := \int_{\frac{i-1}{n}}^{\frac{i}{n}} f(x) dx, \quad (23)$$

for some continuous probability density function $f(x)$. We use the moment function $\phi_j(y) = I(y \leq c_j)$, $j = 1, 2, \dots, m$ where c_j 's are some quantiles. Rather than fixing these c_j , we choose them as the $i/(m+1)$ -quantiles of the N data points. Though introducing a small bias in our procedure, this gives us a reasonable set of quantiles to match against our simulation outputs.

We will use two choices of h , vary the support size n , the number of quantile-based moments m , output data size N , and the function $f(x)$ in our experiments.

Our first choice of $h(\mathbf{X})$ is the time-averaged number of customers in the system during a busy period. Specifically,

$$\begin{aligned} h(X_1, X_2, \dots, X_\tau) &= \frac{\sum_{t=1}^{\tau} (D_t - E_{t-1})}{E\tau}, \\ D_t &= \sum_{k=1}^t X_k, E_t = \sum_{k=1}^t A_k, E_0 = 0, \\ \text{where } \tau &:= \min \left\{ t \geq 1 : \sum_{k=1}^t A_k > \sum_{k=1}^t X_k \right\}. \end{aligned} \quad (24)$$

We set A_t as unit mean exponentials. Here we show that the gradient estimator (12) is valid for this setting. Note that the assumption of finite moment in Proposition 1 Part 2 trivially holds since ϕ_j 's are indicator functions. As for the stopping time τ , if the probability measure governing X_t puts positive weight on each $\frac{d}{n}$, $d = 1, \dots, n$, we have $E[X_t] < 1 = E[A_t]$. It is obvious that $E[e^{\theta(X_t - A_t)}] < \infty$ for some $\theta > 0$. Therefore

$$P(\tau > T) = P\left(\sum_{t=1}^s (X_t - A_t) \geq 0, \forall s \leq T\right) \leq P\left(\sum_{t=1}^T (X_t - A_t) \geq 0\right) \leq e^{-\delta T}, \text{ for all } T \geq 1,$$

for some $\delta > 0$, where the last inequality follows from Lemma 2.6.2 in Durrett (2010). This verifies the assumption on τ because an exponentially decaying tail implies finiteness of moment generating function in a neighborhood of zero. We set $f(x)$ as a uni-modal density

$$f(x) = \frac{1}{\text{Beta}(2, 4)} x (1 - x)^3, \quad (25)$$

and $n = 50$.

We first let the number of observed data $N = 10^5$, a large size so that the empirical moments are very close to the true moments. This is to illustrate the efficacy of our method without introducing another layer of errors due to insufficient output data.

The parameters of Algorithm 1 is set to be $\gamma^k = 10/k$, $M_1 = 300$, $M_2 = 300$. The algorithm is terminated once one of the following conditions is met: (1) the difference between the average of the last 50 iterates and that of the last 51 to 100 iterates is less than 1×10^{-3} , (2) the number of iteration exceeds 1×10^3 . Figure 1 shows a set of typical trace plots with $n = 50$ for different components of \mathbf{p}^k as Algorithm 1 progresses. We can see the evidence of convergence under our stopping criterion and parameter setting.






Figure 1 Trace plots of different components of the input probability vector \mathbf{p} . Support size $n = 50$, number of quantile-based moments $m = 9$. The algorithm terminates after 668 iterations.

Figures 2a and 2b show our input reconstruction results by matching $m = 9$ quantile-based moments. Figure 2a shows the estimated values of $W^*(\eta)$ over a 20-point grid for η between 3.55 and 3.9. We can see that $W^*(\eta)$ stays at zero until around $\eta = 3.7$, and then starts to increase. Thus we identify η^* as approximately 3.7. Note that, even though our MDSA algorithm only guarantees convergence to local optima, the depicted monotone trend of $W^*(\eta)$ shows evidence that our algorithm lands at the global optima. In general, a non-monotone trend alarms that our MDSA algorithm misses global optima, whereas a monotone trend can serve as a validity check of global convergence. Figure 2b shows the optimal \mathbf{p}^* obtained at $\eta^* = 3.7$. The shape of the reconstructed mass function follows the truth quite well, though it is not as smooth. Our method



(a) Optimal value of (4) against η . Run time = 11 min.




(b) $N = 10^5, n = 50, m = 9, \eta^* = 3.713$.


Figure 2 Optimal values of (4) and reconstruction performance in the uni-modal case.

is capable of locating the mode, and correctly gives a decreasing trend towards both sides from the mode.

Next we vary the support size n to investigate its effect on the quality of the estimate yielded by our method. Figures 3a and 3b show the reconstructed distribution, i.e., the estimated $\mathbf{p}^*(\eta^*)$, at $n = 10$ and 100 respectively. The reconstructed distributions both follow quite closely the shape of the truth. The one for $n = 100$ is noticeably more bumpy. This is because n represents essentially the number of parameters in our estimation. As n grows, the estimation variance increases, where the variance comes from both the output data noise and the simulation noise in our reconstruction.



(a) $N = 10^5, n = 10, m = 9, \eta^* = 2.121$, run time = 7 min.



(b) $N = 10^5, n = 100, m = 9, \eta^* = 4.409$, run time = 13 min.

Figure 3 Reconstructed versus the true distribution in the uni-modal case for different n .

Though our analyses and experiments above assume discrete true distributions, we show how our method can be applied in the case of *continuous* true distributions, where we can calibrate

using discrete distributions as approximation. Figure 4 illustrates the reconstructed distributions, using different n and support points $z_k = k/n$, to recover a continuous service time distribution with its density given by (25) based on output data of the time-averaged number of customers in the system during a busy period. To make the comparison fair, we scale down the true density function by a factor of $1/n$ in each of the three cases to obtain the red curves in each of the three plots (this is because the mass at z_k would be approximately $1/n$ times the original density value at z_k , if the true density is discretized over the set of z_k 's). The results are similar to the discrete true distribution setting. In all of the cases $n = 10, 50$, and 100 , the general shape of the truth can be recovered. $n = 10$ has low resolution and cannot capture the sharp shrink of density at the left end. $n = 50$ and $n = 100$ can recover relatively well including near the end points, but $n = 100$ jitters more as the estimation variance grows. In general, there is a tradeoff in choosing an optimal n : A small n leads to a small variance but large bias, whereas a large n leads to a large variance but small bias. In our example one can argue that $n = 50$ makes a balance and is the best among the three cases in capturing the characteristic of the continuous density.






Figure 4 Reconstructed versus the true distribution for different n , for continuous true distribution.

Next we investigate the effect of the number of matched quantiles m . Figures 5a and 5b show the results when we match $m = 4$ and 14 quantiles respectively. While $m = 4$ still appears acceptable, we can see considerable gain at $m = 14$, with the distribution matching almost perfectly for the most part. Compared with $m = 9$ (Figure 2b), $m = 14$ shows slightly better improvement. Here, the data size $N = 10^5$ seems big enough to support matching as many as 14 quantiles without overly introducing variances.

We now test our method in the case of a moderate number of output data points N , to mimic a more realistic situation. Figure 6 shows the reconstructed distribution compared with the truth, for four combinations of $N = 200$ or 500 , and $m = 4$ or 9 . It seems that gathering more data from 200 to 500 does not elicit dramatic improvements. The case of $N = 200$ and $m = 4$ cannot capture the behavior of the left end. However, using $N = 200$, $m = 9$ seems to be able to recover more on that



(a) $N = 10^5, n = 50, m = 4, \eta^* = 3.768$, run time
= 8 min.



(b) $N = 10^5, n = 50, m = 14, \eta^* = 3.694$, run time
= 12 min.

Figure 5 Reconstructed versus the true distribution in the uni-modal case for different m .

end. Of course, we are showing here only with one particular realization of the data, and the details of the reconstruction accuracy could be subject to statistical noise. Nonetheless, it does show that using a relatively large number of quantiles for as small as 200 data can still give promising results.

Next we test our method on a monotone and a bi-modal distribution. Figures 7 and 8 show the results where the true probability mass functions of the service time X_1 are given by (23) with


$$f(x) = \frac{1}{\text{Beta}(\frac{1}{2}, 2)} x^{-\frac{1}{2}} (1-x) \quad \text{and} \quad \frac{0.4}{\text{Beta}(5, 2)} x^4 (1-x) + \frac{0.6}{\text{Beta}(3, 9)} x^2 (1-x)^8,$$

respectively. We use $n = 50, m = 9$ for both settings. Figure 7a shows the values of $W^*(\eta)$ on a grid between $\eta = 2.9$ and 3.9 . The η^* can be approximately located at 3.3 . Figure 7b shows that our reconstructed input distribution recovers the monotonicity of the true mass function very well. On the other hand, Figure 8a and 8b show the results for a bi-modal distribution, in which the reconstruction seems unable to capture the two modes. However, it seems to be capable of capturing the overall trend of more masses on the left than on the right. Figures 9a and 9b also show the similar behaviors of our reconstructions when the true distributions are continuous.


Our second choice of $h(\mathbf{X})$ is the average waiting time of the first 50 customers after starting from an empty single-server queue. In this case,

$$h(X_1, X_2, \dots, X_T) = \frac{1}{T+1} \sum_{t=1}^{T+1} W_t, \quad \text{where } T = 49, W_t = \max\{0, W_{t-1} + X_{t-1} - A_{t-1}\} \text{ for } t \geq 2, W_1 := 0. \quad (26)$$


Note that the validity of gradient estimator (12) is obvious because all ϕ_j 's are indicator functions and τ is a deterministic time. We let A_t follow a known exponential distribution with mean $1/5$. The




(a) $N = 200, n = 50, m = 4, \eta^* = 3.807$, run time = 7 min.



(b) $N = 500, n = 50, m = 4, \eta^* = 3.729$, run time = 8 min.



(c) $N = 200, n = 50, m = 9, \eta^* = 3.768$, run time = 11 min.



(d) $N = 500, n = 50, m = 9, \eta^* = 3.803$, run time = 12 min.

Figure 6 Reconstructed versus the true distribution in the uni-modal case with fewer output data sizes.


support size n and the number of quantile-based moments m are set to be 50 and 9, respectively. The same algorithmic parameter setting and stopping criterion as in the previous set of experiments are adopted. Typical trace plots are given in Figure 10 to demonstrate that the algorithm does converge.

As in the previous set of experiments, our method is tested in the cases of monotone (Figure 11), uni-modal (Figure 12), and multi-modal (Figure 13) distributions, where the underlying probability functions of the service time X_1 are again given by (23) but with $f(x)$ now replaced by


$$\frac{1}{\text{Beta}(2, 0.9)} x (1-x)^{-0.1}, \quad \frac{1}{\text{Beta}(2, 4)} x (1-x)^3,$$

$$\text{and } \frac{0.4}{\text{Beta}(2, 5)} x (1-x)^4 + \frac{0.6}{\text{Beta}(9, 3)} x^8 (1-x)^2,$$

respectively. The results are generally similar to the previous set of experiments. Our method recovers the monotone distribution very well. It cannot fully recover the peaks in the other two




(a) Optimal value of (4) against η . Run time = 9 min.




(b) $N = 10^5, n = 50, m = 9, \eta^* = 3.305$.

Figure 7 Optimal values of (4) and reconstruction performance in the monotone case.



(a) Optimal value of (4) against η . Run time = 11 min.



(b) $N = 10^5, n = 50, m = 9, \eta^* = 3.883$

Figure 8 Optimal values of (4) and reconstruction performance in the multi-modal case.

cases, but it can capture the overall trend of more masses on the left than the right end. Figure 14 reinstates these observations when the true distribution is continuous.

Finally, Figure 15 compares the output probabilities at the matched quantiles from the simulation outputs generated using our reconstructed input model, with those from the output data. We illustrate this comparison for uni-modal, monotone, and bi-modal truths, and for the two types of outputs we have considered. We see that the probabilities are matched very well in all cases. This suggests that for the multi-modal case, the mismatch between our reconstructed distribution and the truth is due to either the insensitivity of the output distribution, or the incapability of the quantiles we use in unveiling the shape of the input distribution.




Figure 9 Continuous true distribution. $n = 50$.




Figure 10 Trace plots of different components of the probability vector \mathbf{p} . Support size $n = 50$, number of quantile-based moments $m = 9$. The algorithm terminates after 619 iterations.

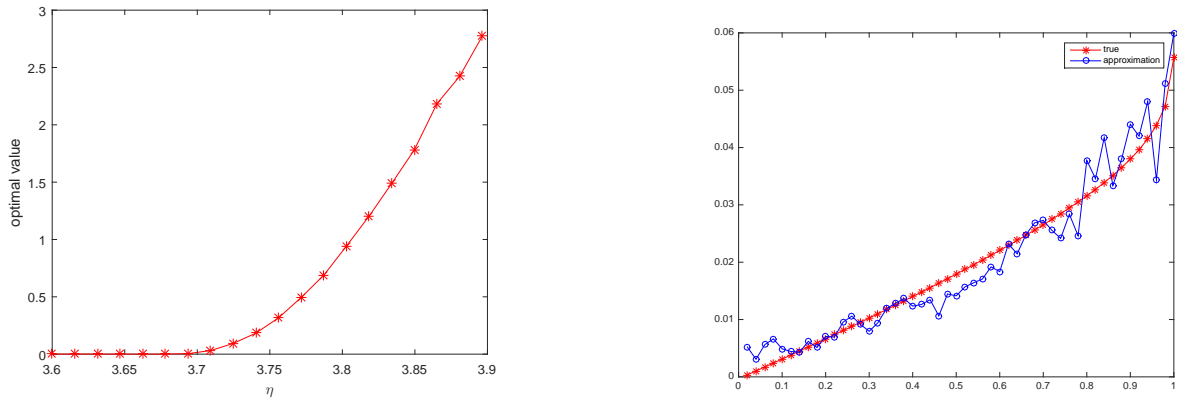




Figure 11 Optimal values of (4) and reconstruction performance in the monotone case.




(a) Optimal value of (4) against η . Run time = 27 min.




(b) $N = 10^5, n = 50, m = 9, \eta^* = 3.725$.

Figure 12 Optimal values of (4) and reconstruction performance in the uni-modal case.




(a) Optimal value of (4) against η . Run time = 26 min.




(b) $N = 10^5, n = 50, m = 9, \eta^* = 3.887$.


Figure 13 Optimal values of (4) and reconstruction performance in the multi-modal case.



(a) Uni-modal case.




(b) Monotone case.




(c) Multi-modal case.


Figure 14 Continuous true distribution. $n = 50$. Y = average wait of first 50 customers starting from an empty system.




(a) Uni-modal case for time-averaged # of customers, i.e., Figure 2b.




(b) Monotone case for time-averaged # of customers, i.e., Figure 7b.




(c) Multi-modal case for time-averaged # of customers, i.e., Figure 8b.



(d) Uni-modal case for 50-customer average wait, i.e., Figure 12b.



(e) Monotone case for 50-customer average wait, i.e., Figure 11b.



(f) Multi-modal case for 50-customer average wait, i.e., Figure 13b.

Figure 15 Moment matches between the simulation output and the true output.

6. Conclusion

We have studied a framework to calibrate the input models in stochastic simulation with only the availability of output data. This inverse model calibration problem appears to be generally understudied in the simulation literature, yet it could arise in many contexts. We have proposed a moment-based approach to nonparametrically infer the input by matching moment statistics, such as quantile probabilities, at the output level. To alleviate the non-identifiability issue, we reduce our model search space by maximizing the entropy among all moment-matching models.

This formulation in general gives a simulation optimization problem that consists of stochastic nonlinear equality constraints. We have converted this difficult problem into a sequence of simulation optimization programs with deterministic convex constraints, and subsequently proposed an MDSA algorithm to solve them efficiently. We have analyzed the convergence properties of our method. Our numerical experiments show that the reconstructed input distributions are generally capable of capturing the overall trends of the truths. They perform particularly well in the case of simple distributions, but with some degradation in the case of more complex distributions.

To enhance the applicability of our methodology, we plan to pursue several future directions. The first is to improve the geometric properties of the reconstructed input model, for instance the

smoothness. The second is methodological development in the case of multiple input and multiple output variables. So far we have only focused on the one-to-one situation, but in practice the simulation models are likely more complex. The third is further investigation on the statistical guarantees and robustness of this line of methods, including also situations where not only the input model is unknown but the system logic in the simulation model could also be subject to errors.

Appendix: Auxiliary Theorems

THEOREM 3 (Corollary in Section 3 in Blum (1954)). *Let Y_k be a sequence of integrable random variables that satisfy*

$$\sum_{k=1}^{\infty} E[E[Y_{k+1} - X_k | X_1, \dots, X_k]^+] < \infty$$

where $x^+ = x$ if $x > 0$ and 0 otherwise, and are bounded below uniformly in k . Then Y_k converges a.s. to a random variable.

LEMMA 2 (Lemma 2.1 in Nemirovski et al. (2009)). *Let ω be defined in (14) and V in (13). Denote \mathcal{X} as the feasible region of (5). For every $\mathbf{q} \in \mathcal{X}$, $\mathbf{p} \in \mathcal{X}^\circ$, and $\boldsymbol{\xi} \in \mathbb{R}^n$, one has*

$$V(\tilde{\mathbf{p}}, \mathbf{q}) \leq V(\mathbf{p}, \mathbf{q}) + \boldsymbol{\xi}'(\mathbf{q} - \mathbf{p}) + \frac{\|\boldsymbol{\xi}\|_*^2}{2\alpha}$$

where $\tilde{\mathbf{p}} = \min_{\mathbf{u} \in \mathcal{X}} \boldsymbol{\xi}'(\mathbf{u} - \mathbf{p}) + V(\mathbf{p}, \mathbf{u})$ is the prox-mapping acting on \mathbf{p} , and $\|\cdot\|_$ is the dual norm of $\|\cdot\|$, α is the strong convexity parameter, both defined in (14).*

Acknowledgments

A preliminary conference version of this work has appeared in Goeva et al. (2014). We gratefully acknowledge support from the National Science Foundation under grants CMMI-1400391/1542020 and CMMI-1436247/1523453. We also thank Peter Haas for suggesting the use of quantile-based moments, and Russell Barton, Shane Henderson and Barry Nelson for many other helpful suggestions.

References

- Avellaneda M, Buff R, Friedman C, Grandechamp N, Kruk L, Newman J (2001) Weighted Monte Carlo: a new technique for calibrating asset-pricing models. *International Journal of Theoretical and Applied Finance* 4(01):91–119.
- Balci O, Sargent RG (1982) Some examples of simulation model validation using hypothesis testing. *Proceedings of the 14th Winter Simulation conference*, volume 2, 621–629 (Winter Simulation Conference).
- Banks J, Carson J, Nelson B, Nicol D (2009) *Discrete-Event System Simulation* (Prentice Hall Englewood Cliffs, NJ, USA), 5th edition edition.

- Barron AR, Sheu CH (1991) Approximation of density functions by sequences of exponential families. *The Annals of Statistics* 19(3):1347–1369.
- Barton RR (2012) Tutorial: Input uncertainty in outout analysis. *Proceedings of the 2012 Winter Simulation Conference (WSC)*, 1–12 (IEEE).
- Barton RR, Nelson BL, Xie W (2013) Quantifying input uncertainty via simulation confidence intervals. *INFORMS Journal on Computing* 26(1):74–87.
- Barton RR, Schruben LW (2001) Resampling methods for input modeling. *Proceedings of the 2001 Winter Simulation Conference*, volume 1, 372–378 (IEEE).
- Basawa I, Bhat U, Zhou J (2008) Parameter estimation using partial information with applications to queueing and related models. *Statistics & Probability Letters* 78(12):1375–1383.
- Basawa IV, Bhat UN, Lund R (1996) Maximum likelihood estimation for single server queues from waiting time data. *Queueing systems* 24(1-4):155–167.
- Beck A, Teboulle M (2003) Mirror descent and nonlinear projected subgradient methods for convex optimization. *Operations Research Letters* 31(3):167–175.
- Benveniste A, Métivier M, Priouret P (2012) *Adaptive Algorithms and Stochastic Approximations*, volume 22 (Springer Science & Business Media).
- Bernardo JM (1979) Expected information as expected utility. *The Annals of Statistics* 686–690.
- Bertsekas DP (1999) *Nonlinear programming* (Athena Scientific).
- Bingham N, Pitts SM (1999) Non-parametric estimation for the $M/G/\infty$ queue. *Annals of the Institute of Statistical Mathematics* 51(1):71–97.
- Blum JR (1954) Multidimensional stochastic approximation methods. *The Annals of Mathematical Statistics* 737–744.
- Box GE, Hill WJ (1967) Discrimination among mechanistic models. *Technometrics* 9(1):57–71.
- Broadie M, Cicek D, Zeevi A (2011) General bounds and finite-time improvement for the Kiefer-Wolfowitz stochastic approximation algorithm. *Operations Research* 59(5):1211–1224.
- Cheng RC, Holland W (1998) Two-point methods for assessing variability in simulation output. *Journal of Statistical Computation Simulation* 60(3):183–205.
- Cheng RC, Holland W (2004) Calculation of confidence intervals for simulation output. *ACM Transactions on Modeling and Computer Simulation (TOMACS)* 14(4):344–362.
- Chick SE (2001) Input distribution selection for simulation experiments: accounting for input uncertainty. *Operations Research* 49(5):744–758.
- Chick SE, Ng SH (2002) Simulation input analysis: joint criterion for factor identification and parameter estimation. *Proceedings of the 34th Winter Simulation Conference*, 400–406 (Winter Simulation Conference).

- Cooper RB (1972) Introduction to queueing theory .
- Cover TM, Thomas JA (1991) Information theory and statistics. *Elements of Information Theory* 279–335.
- Csiszár I (1991) Why least squares and maximum entropy? an axiomatic approach to inference for linear inverse problems. *The Annals of Statistics* 19(4):2032–2066.
- Currin C, Mitchell T, Morris M, Ylvisaker D (1991) Bayesian prediction of deterministic functions, with applications to the design and analysis of computer experiments. *Journal of the American Statistical Association* 86(416):953–963.
- Daley D, Servi L (1998) Moment estimation of customer loss rates from transactional data. *International Journal of Stochastic Analysis* 11(3):301–310.
- Dang CD, Lan G (2015) Stochastic block mirror descent methods for nonsmooth and stochastic optimization. *SIAM Journal on Optimization* 25(2):856–881.
- DeGroot MH (1962) Uncertainty, information, and sequential experiments. *The Annals of Mathematical Statistics* 404–419.
- Donoho DL, Johnstone IM, Hoch JC, Stern AS (1992) Maximum entropy and the nearly black object. *Journal of the Royal Statistical Society. Series B (Methodological)* 41–81.
- Durrett R (2010) *Probability: Theory and Examples* (Cambridge university press).
- Fearnhead P (2004) Filtering recursions for calculating likelihoods for queues based on inter-departure time data. *Statistics and Computing* 14(3):261–266.
- Feng H, Dube P, Zhang L (2014) Estimating life-time distribution by observing population continuously. *Performance Evaluation* 79:182–197.
- Frey JC, Kaplan EH (2010) Queue inference from periodic reporting data. *Operations Research Letters* 38(5):420–426.
- Ghadimi S, Lan G (2013) Stochastic first-and zeroth-order methods for nonconvex stochastic programming. *SIAM Journal on Optimization* 23(4):2341–2368.
- Ghadimi S, Lan G (2015) Accelerated gradient methods for nonconvex nonlinear and stochastic programming. *Mathematical Programming* 1–41.
- Ghosh S, Lam H (2015a) Computing worst-case input models in stochastic simulation. Available at <http://arxiv.org/pdf/1507.05609v1.pdf> .
- Ghosh S, Lam H (2015b) Mirror descent stochastic approximation for computing worst-case stochastic input models. *Proceedings of the 2015 Winter Simulation Conference*, 425–436 (IEEE Press).
- Glasserman P, Yu B (2005) Large sample properties of weighted Monte Carlo estimators. *Operations Research* 53(2):298–312.
- Glynn PW (1990) Likelihood ratio gradient estimation for stochastic systems. *Communications of the ACM* 33(10):75–84.

- Goeva A, Lam H, Zhang B (2014) Reconstructing input models via simulation optimization. *Proceedings of the 2014 Winter Simulation Conference*, 698–709 (IEEE Press).
- Hall AR (2005) *Generalized method of moments* (Oxford University Press).
- Hall P, Park J (2004) Nonparametric inference about service time distribution from indirect measurements. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 66(4):861–875.
- Heckmüller S, Wolfinger BE (2009) Reconstructing arrival processes to G/D/1 queueing systems and tandem networks. *International Symposium on Performance Evaluation of Computer & Telecommunication Systems, 2009. SPECTS 2009.*, volume 41, 361–368 (IEEE).
- Kelton WD, Law AM (2000) *Simulation Modeling and Analysis* (McGraw Hill Boston).
- Kennedy MC, O’Hagan A (2001) Bayesian calibration of computer models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 63(3):425–464.
- Kim YB, Park J (2008) New approaches for inference of unobservable queues. *Proceedings of the 40th Conference on Winter Simulation*, 2820–2825 (Winter Simulation Conference).
- Kleijnen JP (1995) Verification and validation of simulation models. *European Journal of Operational Research* 82(1):145–162.
- Larson RC (1990) The queue inference engine: Deducing queue statistics from transactional data. *Management Science* 36(5):586–601.
- L’Ecuyer P (1990) A unified view of the IPA, SF, and LR gradient estimation techniques. *Management Science* 36(11):1364–1383.
- Lindley DV (1956) On a measure of the information provided by an experiment. *The Annals of Mathematical Statistics* 986–1005.
- Mandelbaum A, Zeltyn S (1998) Estimating characteristics of queueing networks using transactional data. *Queueing systems* 29(1):75–127.
- Marjoram P, Molitor J, Plagnol V, Tavaré S (2003) Markov chain Monte Carlo without likelihoods. *Proceedings of the National Academy of Sciences* 100(26):15324–15328.
- Moulines E, Roueff F, Soulloumiac A, Trigano T (2007) Nonparametric inference of photon energy distribution from indirect measurement. *Bernoulli* 13(2):365–388.
- Nelson B (2016) ‘Some tactical problems in digital simulation’ for the next 10 years. *Journal of Simulation* 10(1):2–11.
- Nemirovski A, Juditsky A, Lan G, Shapiro A (2009) Robust stochastic approximation approach to stochastic programming. *SIAM Journal on Optimization* 19(4):1574–1609.
- Nemirovski A, Yudin D (1983) *Problem Complexity and Method Efficiency in Optimization* (Wiley, New York).

- Park J, Kim YB, Willemain TR (2011) Analysis of an unobservable queue using arrival and departure times. *Computers & Industrial Engineering* 61(3):842–847.
- Pickands III J, Stine RA (1997) Estimation for an $M/G/\infty$ queue with incomplete information. *Biometrika* 84(2):295–308.
- Reiman MI, Weiss A (1989) Sensitivity analysis for simulations via likelihood ratios. *Operations Research* 37(5):830–844.
- Ross JV, Taimre T, Pollett PK (2007) Estimation for queues from queue length data. *Queueing Systems* 55(2):131–138.
- Rubinstein RY (1989) Sensitivity analysis and performance extrapolation for computer simulation models. *Operations Research* 37(1):72–81.
- Santner TJ, Williams BJ, Notz WI (2013) *The Design and Analysis of Computer Experiments* (Springer Science & Business Media).
- Sargent RG (2005) Verification and validation of simulation models. *Proceedings of the 37th Winter Simulation Conference*, 130–143 (Winter Simulation Conference).
- Schruben LW (1980) Establishing the credibility of simulations. *Simulation* 34(3):101–105.
- Serfling RJ (2009) *Approximation Theorems of Mathematical Statistics*, volume 162 (John Wiley & Sons).
- Shirangi MG (2014) History matching production data and uncertainty assessment with an efficient TSVD parameterization algorithm. *Journal of Petroleum Science and Engineering* 113:54–71.
- Song E, Nelson BL, Pegden CD (2014) Advanced tutorial: Input uncertainty quantification. *Proceedings of the 2014 Winter Simulation Conference*, 162–176 (IEEE Press).
- Sra S, Nowozin S, Wright SJ (2012) *Optimization for Machine Learning* (MIT Press).
- Tarantola A (2005) *Inverse Problem Theory and Methods for Model Parameter Estimation* (SIAM).
- Van Campenhout JM, Cover TM (1981) Maximum entropy and conditional probability. *IEEE Transactions on Information Theory* 27(4):483–489.
- Wang TY, Ke JC, Wang KH, Ho SC (2006) Maximum likelihood estimates and confidence intervals of an $M/M/R$ queue with heterogeneous servers. *Mathematical Methods of Operations Research* 63(2):371–384.
- Whitt W (1981) Approximating a point process by a renewal process: The view through a queue, an indirect approach. *Management Science* 27(6):619–636.
- Whitt W (1982) Approximating a point process by a renewal process, I: Two basic methods. *Operations Research* 30(1):125–147.
- Whitt W (2012) Fitting birth-and-death queueing models to data. *Statistics & Probability Letters* 82(5):998–1004.

Wunsch C (1996) *The Ocean Circulation Inverse Problem* (Cambridge University Press).

Zouaoui F, Wilson JR (2004) Accounting for input-model and input-parameter uncertainties in simulation. *IIE Transactions* 36(11):1135–1151.

Supplementary Materials

EC.1. Quadratic Penalty Method

An application of the conventional quadratic penalty method (Bertsekas (1999)) yields the following:

LEMMA EC.1. *Suppose that (3) is feasible. Consider the sequence of optimization programs*

$$\begin{aligned} \min \quad & \sum_{j=1}^m (E_{\mathbf{p}}[\phi_j(h(\mathbf{X}))] - \mu_j)^2 - \lambda R(\mathbf{p}) \\ \text{subject to } & \mathbf{p} \in \mathcal{P} \end{aligned} \quad (\text{EC.1})$$

for $\lambda > 0$. Let $\mathbf{p}^*(\lambda)$ be an optimal solution for (EC.1) indexed at λ . As λ decreases to 0, every limit point of the sequence $\{\mathbf{p}^*(\lambda)\}$ is an optimal solution for (3).

Proof of Lemma EC.1 Consider relaxing the constraints in (3) to get

$$\begin{aligned} \min \quad & -R(\mathbf{p}) + c \sum_{j=1}^m (E_{\mathbf{p}}[\phi_j(h(\mathbf{X}))] - \mu_j)^2 \\ \text{subject to } & \mathbf{p} \in \mathcal{P} \end{aligned} \quad (\text{EC.2})$$

for $c > 0$, which is equivalent to (EC.1) with $\lambda = 1/c$. Proposition 4.2.1 in Bertsekas (1999) entails that as $c \rightarrow \infty$, every limit point of the sequence of optimal solutions for (EC.2) converges to the optimal solution of (3), given that (3) is feasible. This concludes the lemma. \square

EC.2. A Variant of MDSA

In parallel to Section 4.2, we shall design an iterative procedure for solving (EC.1). Note that the objective function in (EC.1) consists of a non-convex, stochastic component $\sum_{j=1}^m (E_{\mathbf{p}}[\phi_j(h(\mathbf{X}))] - \mu_j)^2$ and a convex component $-\lambda R(\mathbf{p})$. We shall use the idea of proximal gradient (Sra et al. (2012)) used for solving composite objective functions in convex problems, which iteratively linearizes the first component while keeping the second component intact at every iteration. The variant of MDSA scheme under this operation amounts to solving, given a current solution \mathbf{p}^k ,

$$\begin{aligned} \min \quad & \gamma^k \hat{\boldsymbol{\psi}}^{k'}(\mathbf{p} - \mathbf{p}^k) - \gamma^k \lambda R(\mathbf{p}) + V(\mathbf{p}^k, \mathbf{p}) \\ \text{subject to } & \mathbf{p} \in \mathcal{P} \end{aligned} \quad (\text{EC.3})$$

The gradient estimate $\hat{\boldsymbol{\psi}}^k$ is the same as discussed in Section 4.2.1, γ^k is the step size, and V is the KL divergence discussed in Section 4.2.2. Consider the generic formulation of (EC.3) written as

$$\begin{aligned} \min \quad & \boldsymbol{\xi}'(\mathbf{q} - \mathbf{p}) - \beta R(\mathbf{q}) + V(\mathbf{p}, \mathbf{q}) \\ \text{subject to } & \mathbf{q} \in \mathcal{P} \end{aligned} \quad (\text{EC.4})$$

LEMMA EC.2. *An optimal solution for (EC.4) is given by $\mathbf{q}^* = (q_1^*, \dots, q_n^*)$ where*

$$q_i^* = \frac{p_i^{\frac{1}{1+\beta}} e^{-\frac{\xi_i}{1+\beta}}}{\sum_{l=1}^n p_l^{\frac{1}{1+\beta}} e^{-\frac{\xi_l}{1+\beta}}} \quad (\text{EC.5})$$

Proof of Lemma EC.2 Consider the Lagrangian for (EC.4)

$$\begin{aligned} & \boldsymbol{\xi}'(\mathbf{q} - \mathbf{p}) - \beta R(\mathbf{q}) + V(\mathbf{p}, \mathbf{q}) + \alpha \left(\sum_{i=1}^n q_i - 1 \right) \\ &= \boldsymbol{\xi}'(\mathbf{q} - \mathbf{p}) + \beta \sum_{i=1}^n q_i \log q_i + \sum_{i=1}^n q_i \log \frac{q_i}{p_i} + \alpha \left(\sum_{i=1}^n q_i - 1 \right) \end{aligned} \quad (\text{EC.6})$$

by relaxing the constraint $\sum_{i=1}^n q_i = 1$. Differentiating (EC.6) with respect to \mathbf{q} gives

$$\xi_i + \beta \log q_i + \beta + \log \frac{q_i}{p_i} + 1 + \alpha$$

Setting to zero gives

$$q_i \propto p_i^{\frac{1}{1+\beta}} e^{-\frac{\xi_i}{1+\beta}}$$

Using the constraint $\sum_{i=1}^n q_i = 1$, we get (EC.5), which can be verified to satisfy the KKT condition straightforwardly. \square

One advantage of using the representation (EC.1) and the stepwise subprogram (EC.3), as compared to (4) and (5) introduced in Section 4, is that it does not involve any root-finding in the MDSA iterations. Thus the resulting procedure is faster than that in Section 4. However, examining when to stop the algorithm becomes less clear as the procedure now relies on the convergence over the sequence of λ to 0, rather than a cutoff at η^* as in Section 4. In the experimental settings in Section 5 we found that it is difficult to determine when to stop in using (EC.1). Since this issue outweighs the marginal advantage in removing the need of solving for a one-dimensional root, we have chosen to adopt (4). Nonetheless, the next section describes that the analysis of the MDSA procedure under this alternate approach in parallel to that under (4).

EC.3. Convergence Analysis of the Variant of MDSA

The variant of MDSA for solving (EC.1) is depicted in Algorithm 2. Similar to Algorithm 1, Steps 2 and 3 in Algorithm 1 combine to solve (EC.3) with \mathcal{P} replaced by $\mathcal{P}(\epsilon)$. The rationale for such operations lies in a technicality in guaranteeing boundedness of the gradient estimator, and subsequently algorithmic convergence, as in Algorithm 1.

To prove almost sure convergence of Algorithm 2, we need the following generalization of Lemma 2.1 in Nemirovski et al. (2009):

LEMMA EC.3. *Let \mathcal{X} be a convex set in \mathbb{R}^n and $\|\cdot\|$ be a norm with dual $\|\cdot\|_*$. Let $\omega : \mathcal{X} \rightarrow \mathbb{R}$ be a strongly convex function that satisfies*

$$\omega(z) \geq \omega(x) + \nabla \omega(x)'(z - x) + \frac{\alpha}{2} \|z - x\|^2$$

Algorithm 2 MDSA for solving (4)

Input: A small parameter $\epsilon > 0$, initial solution $\mathbf{p}^1 \in \mathcal{P}(\epsilon) = \{\mathbf{p} : \sum_{i=1}^n p_i = 1, p_i \geq \epsilon \text{ for } i = 1, \dots, n\}$, a step size sequence γ^k , and sample sizes M_1 and M_2 .

Iteration: For $k = 1, 2, \dots$, do the following: Given \mathbf{p}^k ,

1. Estimate $\hat{\psi}^k = (\hat{\psi}_1^k, \dots, \hat{\psi}_n^k)$ with

$$\hat{\psi}_i^k = 2 \sum_{j=1}^m \frac{1}{M_1} \sum_{r=1}^{M_1} (\phi_j(h(\mathbf{X}^{(r)})) - \mu_j) \frac{1}{M_2} \sum_{r=1}^{M_2} \phi_j(h(\tilde{\mathbf{X}}^{(r)})) S_i(\tilde{\mathbf{X}}^{(r)}; \mathbf{p}^k)$$

where $\mathbf{X}^{(r)}$ and $\tilde{\mathbf{X}}^{(r)}$ are M_1 and M_2 independent copies of the input process generated under i.i.d. replications of \mathbf{p}^k , which are used simultaneously for all components of $\hat{\psi}^k$.

2. Output

$$p_i^{k+1} = \frac{p_i^k \frac{1}{1+\gamma^k \lambda} e^{-\frac{\gamma^k \hat{\psi}_i^k}{1+\gamma^k \lambda}}}{\sum_{l=1}^n p_l^k \frac{1}{1+\gamma^k \lambda} e^{-\frac{\gamma^k \hat{\psi}_l^k}{1+\gamma^k \lambda}}}$$

3. If $p_i^{k+1} < \epsilon$ for some i , then solve the convex optimization (EC.3) but with \mathcal{P} replaced by the set $\mathcal{P}(\epsilon)$. Output its solution as \mathbf{p}^{k+1} .

for any $x, z \in \mathcal{X}$. Let $f : \mathcal{X} \rightarrow \mathbb{R}$ be a convex differentiable function. Define, for any $x, z \in \mathcal{X}$,

$$V_1(x, z) = \omega(z) - \omega(x) - \nabla \omega(x)'(z - x)$$

$$V_2(x, z) = f(z) - f(x) - \nabla f(x)'(z - x)$$

and $\tilde{V}(x, z) = V_1(x, z) + V_2(x, z)$. Then, given any $x, y \in \mathbb{R}^n$, we have

$$\tilde{V}(v, u) - \tilde{V}(x, u) \leq f(x) - f(v) + (y + \nabla f(x))'(u - x) + \frac{\|y\|_*^2}{2\alpha}$$

for

$$v = \operatorname{argmin}_{z \in \mathcal{X}} \{y'(z - x) + f(z) + V_1(x, z)\} \quad (\text{EC.7})$$

and any $u \in \mathcal{X}$.

Proof of Lemma EC.3 Given $x, y \in \mathbb{R}^n$, define v as in (EC.7), and consider

$$\begin{aligned} & V_1(v, u) - V_1(x, u) \\ &= \omega(u) - \omega(v) - \nabla \omega(v)'(u - v) - (\omega(u) - \omega(x) - \nabla \omega(x)'(u - x)) \\ &= \omega(x) - \omega(v) - \nabla \omega(v)'(u - v) + \nabla \omega(x)'(u - x) \\ &= (\nabla \omega(x) - \nabla \omega(v) - \nabla f(v) - y)'(u - v) + \nabla \omega(x)'(v - x) + y'(u - v) + \nabla f(v)'(u - v) + \omega(x) - \omega(v) \\ &\leq y'(u - v) + \nabla f(v)'(u - v) - V_1(x, v) \end{aligned} \quad (\text{EC.8})$$

where the inequality follows from $(y + \nabla f(v) + \nabla \omega(v) - \nabla \omega(x))'(u - v) \geq 0$, by the optimality of v on the convex function $y'(z - x) + f(z) + V(x, z)$ in z . On the other hand,

$$V_2(v, u) - V_2(x, u) = f(x) - f(v) - \nabla f(v)'(u - v) + \nabla f(x)'(u - x) \quad (\text{EC.9})$$

Hence, from (EC.8) and (EC.9), we have

$$\begin{aligned} & \tilde{V}(v, u) - \tilde{V}(x, u) \\ & \leq y'(u - v) + f(x) - f(v) + \nabla f(x)'(u - x) - V_1(x, v) \\ & = y'(x - v) + f(x) - f(v) + (y + \nabla f(x))'(u - x) - V_1(x, v) \end{aligned} \quad (\text{EC.10})$$

Using $y'(x - v) \leq \frac{\|y\|_*^2}{2\alpha} + \frac{\alpha}{2}\|x - v\|^2$, via Young's inequality (Nemirovski et al. (2009)), and $V_1(x, v) \geq \frac{\alpha}{2}\|x - v\|^2$ from the definition of V_1 , we have (EC.10) less than or equal to

$$f(x) - f(v) + (y + \nabla f(x))'(u - x) + \frac{\|y\|_*^2}{2\alpha}$$

□

The following is the analog of Theorem 2 for Algorithm 2:

THEOREM EC.1. *Suppose there exists a unique optimal solution $\mathbf{p}^* \in \mathcal{P}(\epsilon)$ for (4) such that $(\psi(\mathbf{p}) - \lambda \nabla R(\mathbf{p}))'(\mathbf{p} - \mathbf{p}^*) = 0$ if and only if $\mathbf{p} = \mathbf{p}^*$. Choose a non-increasing step size sequence $\{\gamma^k\}$ such that*

$$\sum_{k=1}^{\infty} \gamma^k = \infty, \quad \sum_{k=1}^{\infty} \gamma^{k^2} < \infty$$

Then \mathbf{p}^k generated in Algorithm 2 converges to \mathbf{p}^ .*

Note that $\psi(\mathbf{p}) - \lambda \nabla R(\mathbf{p})$ is now the gradient (with the first term obtained from a perturbation within the probability simplex) of the corresponding objective function in (EC.1), and $(\psi(\mathbf{p}) - \lambda \nabla R(\mathbf{p}))'(\mathbf{p} - \mathbf{p}^*) \geq 0$ by the optimality of \mathbf{p}^* . We also mention that the condition of non-increasing $\{\gamma^k\}$ can be replaced readily by eventually non-increasing $\{\gamma^k\}$.

Proof of Theorem EC.1 By defining $\mathcal{X} = \mathcal{P}$, $\omega(\mathbf{p}^k) = \sum_{i=1}^n p_i^k \log p_i^k$, $f(\mathbf{p}) = -\gamma^k \lambda R(\mathbf{p}^k) = \gamma^k \lambda \sum_{i=1}^n p_i^k \log p_i^k$, $x = \mathbf{p}^k$, $y = \gamma^k \hat{\psi}^k$, $u = \mathbf{p}^*$, $\|\cdot\|$ as L_1 -norm and $\|\cdot\|_*$ as the supremum norm in Lemma EC.3, we have $\tilde{V}(\mathbf{p}^k, \mathbf{p}^*) = (1 + \gamma^k \lambda) \sum_{i=1}^n p_i^* \log \frac{p_i^*}{p_i^k}$, $\alpha = 1$, and

$$\tilde{V}(\mathbf{p}^{k+1}, \mathbf{p}^*) - \tilde{V}(\mathbf{p}^k, \mathbf{p}^*) \leq \gamma^k \lambda (R(\mathbf{p}^{k+1}) - R(\mathbf{p}^k)) + \gamma^k (\hat{\psi}^k - \lambda \nabla R(\mathbf{p}^k))'(\mathbf{p}^* - \mathbf{p}^k) + \frac{\gamma^{k^2} \|\hat{\psi}^k\|_\infty^2}{2} \quad (\text{EC.11})$$

Let $V(\mathbf{p}, \mathbf{q}) = \sum_{i=1}^n q_i^k \log \frac{q_i^k}{p_i^k}$ be the KL divergence, so that $\tilde{V}(\mathbf{p}^k, \mathbf{p}^*) = (1 + \gamma^k \lambda) V(\mathbf{p}^k, \mathbf{p}^*)$. Let \mathcal{F}^k be the filtration generated by $\{\mathbf{p}^1, \dots, \mathbf{p}^k\}$. Taking conditional expectation on (EC.11), we have

$$\begin{aligned} & E[V(\mathbf{p}^{k+1}, \mathbf{p}^*) - V(\mathbf{p}^k, \mathbf{p}^*) | \mathcal{F}^k] \\ & \leq \lambda \gamma^k (V(\mathbf{p}^k, \mathbf{p}^*) - E[V(\mathbf{p}^{k+1}, \mathbf{p}^*) | \mathcal{F}^k]) - \lambda \gamma^k (R(\mathbf{p}^k) - E[R(\mathbf{p}^{k+1}) | \mathcal{F}^k]) + \gamma^k (\psi(\mathbf{p}^k) - \lambda \nabla R(\mathbf{p}^k))' (\mathbf{p}^* - \mathbf{p}^k) \\ & \quad + \frac{\gamma^{k^2} E[\|\hat{\psi}^k\|_\infty^2 | \mathcal{F}^k]}{2} \end{aligned} \tag{EC.12}$$

Taking expectation and summing up on both sides of (EC.12), we have

$$\begin{aligned} & \sum_{k=1}^K E[E[V(\mathbf{p}^{k+1}, \mathbf{p}^*) - V(\mathbf{p}^k, \mathbf{p}^*) | \mathcal{F}^k]^+] \\ & \leq \lambda \sum_{k=1}^K \gamma^k (E[V(\mathbf{p}^{k+1}, \mathbf{p}^*)] - E[V(\mathbf{p}^k, \mathbf{p}^*)]) - \lambda \sum_{k=1}^K \gamma^k (E[R(\mathbf{p}^k)] - E[R(\mathbf{p}^{k+1})]) \\ & \quad + \sum_{k=1}^K \gamma^k E[(\psi(\mathbf{p}^k) - \lambda \nabla R(\mathbf{p}^k))' (\mathbf{p}^* - \mathbf{p}^k)] + \sum_{k=1}^K \frac{\gamma^{k^2} E\|\hat{\psi}^k\|_\infty^2}{2} \\ & = \lambda \left(\gamma^1 V(\mathbf{p}^1, \mathbf{p}^*) + \sum_{k=2}^K (\gamma^k - \gamma^{k-1}) E[V(\mathbf{p}^k, \mathbf{p}^*)] - \gamma^K E[V(\mathbf{p}^{K+1}, \mathbf{p}^*)] \right) - \lambda \left(\gamma^1 R(\mathbf{p}^1) \right. \\ & \quad \left. + \sum_{k=2}^K (\gamma^k - \gamma^{k-1}) E[R(\mathbf{p}^k, \mathbf{p}^*)] - \gamma^K E[R(\mathbf{p}^{K+1}, \mathbf{p}^*)] \right) + \sum_{k=1}^K \gamma^k E[(\psi(\mathbf{p}^k) - \lambda \nabla R(\mathbf{p}^k))' (\mathbf{p}^* - \mathbf{p}^k)] \\ & \quad + \sum_{k=1}^K \frac{\gamma^{k^2} E\|\hat{\psi}^k\|_\infty^2}{2} \end{aligned} \tag{EC.13}$$

by telescoping. Note that $V(\mathbf{p}, \mathbf{p}^*) \geq 0$, $R(\mathbf{p}) \geq 0$, and $R(\mathbf{p}) \leq \log n$ for any $\mathbf{p} \in \mathcal{P}$. Moreover, $(\psi(\mathbf{p}) - \lambda \nabla R(\mathbf{p}^k))' (\mathbf{p} - \mathbf{p}^*) \geq 0$ for any $\mathbf{p} \in \mathcal{P}$. Also, by the same argument as in the proof of Theorem 2, we have $E[\|\hat{\psi}^k\|_\infty] \leq C$ uniformly for some $C > 0$. Therefore, (EC.13) is less than or equal to

$$\lambda \gamma^1 V(\mathbf{p}^1, \mathbf{p}^*) + \lambda \sum_{k=2}^K (\gamma^{k-1} - \gamma^k) \log n + \lambda \gamma^K \log n + \frac{C^2}{2} \sum_{k=1}^K \gamma^{k^2} = \lambda \gamma^1 V(\mathbf{p}^1, \mathbf{p}^*) + \lambda \gamma^1 \log n + \frac{C^2}{2} \sum_{k=1}^K \gamma^{k^2}$$

by telescoping. Letting $K \rightarrow \infty$, we have

$$\sum_{k=1}^{\infty} E[E[V(\mathbf{p}^{k+1}, \mathbf{p}^*) - V(\mathbf{p}^k, \mathbf{p}^*) | \mathcal{F}^k]^+] \leq \lambda \gamma^1 V(\mathbf{p}^1, \mathbf{p}^*) + \lambda \gamma^1 \log n + \frac{C^2}{2} \sum_{k=1}^{\infty} \gamma^{k^2} < \infty$$

By martingale convergence theorem (Theorem 3 in the Appendix), we have $V(\mathbf{p}^k, \mathbf{p}^*)$ converges a.s. to some integrable random variable V_∞ .

Now, taking expectation and summing up on both sides of (EC.11), and by a similar argument as above, we have

$$\sum_{k=1}^{\infty} \gamma^k E[\psi(\mathbf{p}^k)' (\mathbf{p}^k - \mathbf{p}^*)] \leq V(\mathbf{p}^1, \mathbf{p}^*) - EV_\infty + \lambda \gamma^1 V(\mathbf{p}^1, \mathbf{p}^*) + \lambda \gamma^1 \log n + \frac{C^2}{2} \sum_{k=1}^{\infty} \gamma^{k^2} < \infty$$

As in the proof of Theorem 2, since $\sum_{k=1}^{\infty} \gamma^k = \infty$ and $E[(\psi(\mathbf{p}^k) - \lambda \nabla R(\mathbf{p}^k))'(\mathbf{p}^k - \mathbf{p}^*)] \geq 0$, for all k , there must exist a subsequence k_i such that $E[(\psi(\mathbf{p}^{k_i}) - \lambda \nabla R(\mathbf{p}^{k_i}))'(\mathbf{p}^{k_i} - \mathbf{p}^*)] \rightarrow 0$ a.s.. This implies that $(\psi(\mathbf{p}^{k_i}) - \lambda \nabla R(\mathbf{p}^{k_i}))'(\mathbf{p}^{k_i} - \mathbf{p}^*) \xrightarrow{P} 0$, which in turn implies the existence of a further subsequence l_i such that $(\psi(\mathbf{p}^{l_i}) - \lambda \nabla R(\mathbf{p}^{l_i}))'(\mathbf{p}^{l_i} - \mathbf{p}^*) \rightarrow 0$ a.s.. From Proposition 1 Part 2, we have $\psi(\mathbf{p})$ continuous in \mathbf{p} . By the assumption that $(\psi(\mathbf{p}) - \lambda \nabla R(\mathbf{p}))'(\mathbf{p} - \mathbf{p}^*) = 0$ only if $\mathbf{p} = \mathbf{p}^*$, and that $(\psi(\mathbf{p}) - \lambda \nabla R(\mathbf{p}))'(\mathbf{p} - \mathbf{p}^*)$ is continuous in \mathbf{p} , we have $\mathbf{p}^{l_i} \rightarrow \mathbf{p}^*$ a.s.. Hence $V(\mathbf{p}^{l_i}, \mathbf{p}^*) \rightarrow 0$ a.s.. Since we have proved above that $V(\mathbf{p}^k, \mathbf{p}^*)$ converges a.s., this limit must be 0. Therefore, by Pinsker's inequality, we have $\mathbf{p}^k \rightarrow \mathbf{p}^*$ in total variation a.s.. This concludes the theorem. \square